

**An efficient algorithm for sensitivity analysis of
chaotic systems**

by

Nisha Chandramoorthy

Submitted to the Department of Mechanical Engineering and Center for
Computational Science and Engineering

in partial fulfillment of the requirements for the degree of

Ph.D. in Mechanical Engineering and Computation

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2021

© Massachusetts Institute of Technology 2021. All rights reserved.

Author
Department of Mechanical Engineering and Center for Computational
Science and Engineering
31st July, 2021

Certified by.....
Qiqi Wang
Associate Professor, Department of Aeronautics and Astronautics
Thesis Supervisor

Accepted by
Nicolas Hadjiconstantinou
Professor, Department of Mechanical Engineering
Chairman, Department Committee on Graduate Theses

Accepted by
Youssef Marzouk
Professor, Department of Aeronautics and Astronautics
Co-director, Center for Computational Science and Engineering

An efficient algorithm for sensitivity analysis of chaotic systems

by

Nisha Chandramoorthy

Submitted to the Department of Mechanical Engineering and Center for Computational Science and Engineering
on 31st July, 2021, in partial fulfillment of the
requirements for the degree of
Ph.D. in Mechanical Engineering and Computation

Abstract

How does long-term chaotic behavior respond to small parameter perturbations? Using detailed models, chaotic systems are frequently simulated across disciplines – from climate science to astrophysics. But, an efficient computation of parametric derivatives of their statistics or long-term averages, also known as linear response, is an open problem. The difficulty is due to an inherent feature of chaos: an exponential growth over time of infinitesimal perturbations, which renders conventional methods for sensitivity computation inapplicable. More sophisticated recent approaches, including ensemble-based and shadowing-based methods are either computationally impractical or lack convergence guarantees. We propose a novel alternative known as space-split sensitivity or S3, which evaluates linear response as an efficiently computable, provably convergent ergodic average. The main contribution of this thesis is the development of the S3 algorithm for uniformly hyperbolic systems – the simplest setting in which chaotic attractors occur – with one-dimensional unstable manifolds. S3 can enable applications of the computed sensitivities to optimization, control theory and uncertainty quantification, in the realm of chaotic dynamics, wherein these applications remain nascent.

We propose a transformation of Ruelle’s rigorous linear response formula, which is ill-conditioned in its original form, into a well-conditioned ergodic-averaging computation. We prove a decomposition of Ruelle’s formula, called the S3 decomposition, that is differentiable on the unstable manifold. The S3 decomposition ensures that one of the resulting terms, the stable contribution, can be computed using a regularized tangent equation, similar to in a non-chaotic system. The remainder, known as the unstable contribution, is regularized and converted into a computable ergodic average. The S3 algorithm presented here can be naturally extended to systems with higher-dimensional unstable manifolds.

The secondary contributions of this thesis are analysis and applications of existing methods, including those shadowing-based and ensemble-based, to compute linear response. A feasibility analysis of ensemble sensitivity calculation, which is a direct

evaluation of Ruelle's formula, reveals a problem-dependent, typically poor rate of convergence, rendering it computationally impractical. Shadowing-based sensitivity computation is not guaranteed to converge because of atypicality of shadowing orbits. This atypicality also implies that small parameter perturbations can lead, contrary to popular belief, to a large change in the statistics of a chaotic system, a consequence being that numerical simulations of chaotic systems may not reproduce their true long-term behaviors.

Thesis Supervisor: Qiqi Wang

Title: Associate Professor, Department of Aeronautics and Astronautics

Acknowledgments

I feel lucky to have been advised by Prof. Qiqi Wang. He has been a generous teacher and a brilliant collaborator who has made every result in this thesis possible, and helped me develop as a researcher. I am greatly indebted to my mentors and thesis committee members Prof. Semyon Dyatlov and Prof. Youssef Marzouk for their insight and advice at just the right times. I am grateful to my thesis committee members Prof. Pierre Lermusiaux and Prof. Themis Sapsis for their valuable time and suggestions, which have improved this work. The analysis of ensemble sensitivity computation (Chapter 4) was done jointly with my collaborators and friends, Pablo Fernandez and Chaitanya Talnikar. I am grateful to have worked with Prof. Luca Magri on applications of automatic differentiated-shadowing (Chapter 2). Initial work for this project began during my internship with Dr. Sri Hari Krishna Narayanan and Dr. Paul Hovland, who were excellent mentors and collaborators. I am extremely grateful to Qiqi for sharing his results on the atypicality of shadowing orbits (Chapter 3). I am thankful to Angxiu Ni for pointing out the limiting assumption in the pre-S3 algorithm (Chapter 5), which led us to develop S3 (Chapter 7). My sincere thanks to Malo Jezequel for formally proving the expression for the stable contribution (section 7.7), and teaching me a lot in the process.

In the long development of this thesis work, I have cultivated profound respect and gratitude for many mathematicians and fluid dynamicists, who have expended time and effort to share their knowledge with me, including Qiqi Wang, Vaughn Climenhaga, Semyon Dyatlov, Malo Jezequel, Kevin Lin, Peter Schmid, George Rigas, Carlangelo Liverani, Gary Froyland, Johan Larsson and David Ruelle (who pioneered the study of linear response). I am thankful to my mentors, including Prof. Mary Catherine Kropinski, Dr. Scott Murman, Dr. Marie Denison, Prof. Praveen Chandrashekar, Prof. Nilima Nigam and Prof. Nicolas Hadjiconstantinou, whom I am extremely fortunate to have met during my undergraduate and graduate student years. I also thank Kate Nelson, Jean Sofronas, Leslie Regan and Saana McDaniel for all their kindness and timely administrative support.

At times, my journey has resembled eating plentiful meals and resting peacefully in the Shire much less or not at all, when compared to wading through a labyrinthine Shelob's lair without a sword to protect myself. But, some people have made the wading seem shorter and easier by taking on many avatars as fellow runners, ice-cream eaters, non-judgmental appraisers of strange obsessions and doomsday theories, receptacles for long rants, stoic program debuggers and so on. For all that and going beyond, I am forever grateful to Karthik, Jerry, Rachel, Sneha, Max (who also suggested the epigraph in the introduction), Ben Z, Priya, Yawar, Essie, Cory, Akshay, Sarah, Philip, Chandana, Lee, Rohit, Saviz, Ben C, Hugh, Aarti, Ricardo, I-Ting, Boyu, Cameron, Nigamaa, Rashmi, Matt, Shibani and Shun. I have the privilege of having an extended family that cares about me. Thank you for your concern and friendship, Karthik, Chitti, Lekha, Saranya, Chittappa, Pinma, Harini, Deepak, Swathi, Raja, Prakash, Abhi, Ashwini, Avvas and Thatha. My eternal gratitude goes to my parents, Hemalatha Sundaresan and V S Chandramoorthy, for their love and unwavering support. Finally, thank you, Nandhini Chandramoorthy, always the better sister, for everything.

Contents

1	Introduction	27
1.1	Long-term or statistical response to infinitesimal perturbations and its applications	27
1.2	Existing methods for linear response computation	31
1.3	Thesis contributions	32
1.4	Outline	33
2	Shadowing and its application to a time-delayed acoustic model	35
2.1	Introduction	36
2.2	Shadowing sensitivity in chaotic systems	39
2.2.1	Tangent dynamics	41
2.2.2	Adjoint dynamics	43
2.2.3	Automatic differentiation to compute tangent and adjoint solutions	45
2.2.4	The problem with computing sensitivities of ergodic averages using conventional methods	47
2.2.5	Non-Intrusive Least Squares Shadowing	49
2.2.6	Error in shadowing and alternatives	53
2.3	Tangent/adjoint shadowing algorithm	54
2.3.1	Shadowing algorithm: inputs and outputs	54
2.3.2	Evolution of homogeneous and inhomogeneous perturbations with repeated normalization	55
2.3.3	Minimizing the growth of the shadowing perturbation sequence	57

2.3.4	Modifications due to the center direction	59
2.3.5	Computation of the sensitivities	60
2.3.6	Automatic differentiation	61
2.4	The time delayed model for thermoacoustics	61
2.4.1	Replacing the time delay with an advection equation	63
2.4.2	Types of solutions over a range of the heat-release parameter	65
2.4.3	Acoustic energy and Rayleigh criterion as objective functions	68
2.4.4	Lyapunov vectors in the chaotic regime	71
2.5	Suppression of a nonlinear oscillation by gradient-based optimization	72
2.5.1	Shadowing sensitivities of the acoustic energy and Rayleigh criterion	73
2.5.2	Minimization of the acoustic energy using shadowing	76
2.6	Data assimilation with discrete shadowing	78
2.6.1	Tangent NILSS for state estimation with full-state observations	80
2.6.2	Converting the state estimation problem to a parameter estimation problem	81
2.6.3	Numerical results	83
2.7	Conclusions	84
2.8	Validation of sensitivity computation on the Lorenz'63 model	85
2.9	Validation of data assimilation scheme on the Lorenz'63 model	86
2.10	AD shadowing and code for replication/extension	87
3	On the probability of finding a nonphysical solution through shadowing	91
3.1	Introduction	91
3.2	Shadowing and applications	93
3.3	What are physical and nonphysical solutions?	94
3.3.1	Physical solutions	95
3.3.2	Nonphysical solutions type I: Periodic solutions	99
3.3.3	Tent map: periodic and physical solutions	102

3.3.4	Nonphysical solutions type II: Quasi-physical solutions	104
3.4	Are shadowing solutions physical?	108
3.4.1	Example of a shadowing solution for the tent map	108
3.4.2	An example of quasi-physical shadowing solution	110
3.4.3	Are general shadowing solutions physical?	113
3.4.4	What about the Lorenz equation?	119
3.5	Discussion and outlook	123
3.6	Supplementary material	126
3.6.1	Approximation of the Lorenz Map	126
3.6.2	Computing shadowing solutions	128
3.6.3	The <i>plucked</i> tent map	129
4	Ruelle’s formula and the infeasibility of its direct evaluation	133
4.1	Introduction	133
4.2	The ES estimator	136
4.2.1	The sensitivity computation problem setup	136
4.2.2	The Lea-Allen-Haine ES algorithm	137
4.3	Error vs. computational cost of the ES estimator	139
4.3.1	ES estimator as approximation of Ruelle’s response formula	140
4.3.2	Bias and variance of the ES estimator	142
4.3.3	Optimistic convergence estimate for the ES method	146
4.4	Numerical examples	148
4.4.1	The Lorenz ’63 attractor	149
4.4.2	The Lorenz ’96 model	153
4.4.3	Chaotic flow over an airfoil	155
4.4.4	Turbulent flow over a turbine vane	159
4.5	Discussion and comments	162
4.6	Conclusion	165
5	Computing linear response to unstable perturbations	167
5.1	Introduction	168

5.2	Problem statement	170
5.2.1	The primal dynamics	171
5.2.2	Ensemble and ergodic averages	171
5.2.3	Quantity of interest	172
5.2.4	Ruelle’s formula and its computational inefficiency	172
5.3	Main contributions	174
5.4	Derivation of the pre-S3 formula	176
5.4.1	Ruelle’s formula split along Oseledets spaces	177
5.4.2	Derivation of the stable contribution	178
5.4.3	Computation of the stable contribution	178
5.4.4	The unstable contribution: an ansatz	179
5.4.5	Computation of the unstable contribution	184
5.5	Interpretation of the unstable contribution	185
5.6	Comments on pre-S3 computation	188
5.7	Numerical examples	190
5.7.1	Smale-Williams solenoid map	190
5.7.2	Kuznetsov-Plykin map	191
5.8	Conclusions	192
5.9	Appendix	193
6	Computing the curvature of one-dimensional unstable manifolds	199
6.1	Introduction	200
6.2	Problem setup, definitions and review of Covariant Lyapunov Vectors	202
6.2.1	Tangent dynamics	203
6.2.2	Uniform hyperbolicity	205
6.2.3	Examples	205
6.2.4	Lack of differentiability of E^u and E^s	206
6.2.5	Derivatives of CLVs in their own directions	207
6.2.6	Computations along trajectories	208

6.3	An algorithm to compute the directional derivatives of CLVs in their own directions	208
6.3.1	Change of coordinates and associated notation	209
6.3.2	The map in adapted coordinates	210
6.3.3	Computation of unstable CLV self-derivatives	211
6.3.4	The differential CLV method: iterative orthogonal projections	212
6.4	Numerical results implementing the differential CLV method	214
6.4.1	Validation against analytical curvature of the Solenoid map . .	214
6.4.2	Numerical verification of the curvature of the Lorenz attractor	217
6.4.3	Qualitative verification on a perturbed cat map	222
6.4.4	Qualitative verification on the volume-decreasing perturbed Cat	224
6.4.5	Numerical results on the Hènon map	228
6.5	An application of CLV derivatives to statistical linear response	230
6.6	Conclusion	232
6.7	Appendix	233
6.7.1	The lack of differentiability of CLVs	233
6.7.2	Computations on the super-contracting Solenoid attractor . .	234
6.8	Convergence of the differential CLV method	235
6.8.1	Regularization of Ruelle’s formula	235
7	Space-split sensitivity algorithm for one-dimensional unstable manifolds	239
7.1	Introduction	240
7.2	Preliminaries and problem setup	244
7.2.1	Ergodic theory and linear response	244
7.2.2	Tangent dynamics	245
7.2.3	Chaotic systems	246
7.2.4	Uniform hyperbolicity	247
7.2.5	The SRB measure	248
7.2.6	Parameterization of unstable manifolds	249

7.2.7	Iterative differentiation on the unstable manifold	249
7.2.8	Conditional density of the SRB measure on unstable manifolds	251
7.2.9	Linear response of uniformly hyperbolic systems	252
7.3	Main results: S3 decomposition and computation of Ruelle's response formula	253
7.3.1	The S3 algorithm	255
7.4	Derivation and computation of the S3 formula	257
7.4.1	Regularizing tangent equation solutions	258
7.4.2	Alternative expression of the stable contribution	259
7.4.3	Computation of the stable contribution	260
7.4.4	Alternative expression of the unstable contribution	260
7.5	Computing derivatives along unstable manifolds	262
7.5.1	Iteratively computing the curvature of the unstable manifold .	263
7.5.2	Iterative formula for unstable derivatives of SRB density . . .	265
7.5.3	Iterative computation of the scalar field b	266
7.5.4	Computation of the unstable contribution	268
7.6	Numerical results on perturbations of Baker's map	271
7.6.1	Stable and unstable subspaces	272
7.6.2	SRB measures of perturbed Baker's maps	272
7.6.3	Sensitivities of a smooth objective function	276
7.7	Proof of Theorem 1	277
7.7.1	Existence and uniqueness of a regularized perturbation field .	279
7.7.2	Existence of self-derivative of the unstable direction	281
7.7.3	Convergence of the unstable projections of the regularized tan- gent solutions	283
7.7.4	Route to showing existence and differentiability of the S3 de- composition	284
7.7.5	Differentiability of v in the unstable direction	285
7.7.6	Convergence of the derivative of $\{a^n\}$	289
7.8	Proof of Theorem 2	290

7.8.1	Convergence of the stable contribution	290
7.8.2	Convergence of the unstable contribution	292
7.9	Discussion and conclusion	295
8	Extending S3 to Chaotic ODEs	297
8.1	Chaotic ODE	298
8.2	Uniformly hyperbolic flows	298
8.2.1	The center subspace	298
8.3	Center-Unstable projection	299
8.4	Stable-Unstable-Center decomposition	301
8.5	Center contribution alternative expression	302
8.6	Updating the unstable contribution	303
8.6.1	The existence of unstable derivatives of a and c	303
8.7	Differentiating the unstable projection	306
8.8	The S3 algorithm for uniformly hyperbolic flows	308
8.9	Numerical example: Lorenz '63 system	311
8.10	Conclusion	313
9	Remarks on the S3 algorithm	315
9.1	Perturbations of the Solenoid Map	315
9.2	Finite sample error analysis of the S3 algorithm	317
9.3	Finite sample cost analysis of S3	322
9.3.1	Discrete adjoint of the S3 algorithm	324
10	Summary and Future Work	327
10.1	S3 extension to multi-dimensional unstable manifolds	328

List of Figures

1-1	Linear perturbation solutions computed for the thermoacoustic system discussed in Chapter 2. FD stands for finite-difference and AD for automatic differentiation.	29
2-1	l^2 norms of the perturbation vectors computed through the homogeneous tangent (blue triangle), adjoint equations (orange circle), finite difference (green Y), forward-mode AD (red triangle) and reverse-mode AD (purple plus), are shown as a function of time, for the time-delayed model of section 2.4.	47
2-2	Bifurcation diagram. The values of $\langle J_{ac} \rangle$, the time-averaged acoustic energy, are color-coded according to the type of solution: periodic (blue), quasiperiodic (green) and chaotic (red). The attractors represented on the u_f - \dot{q} plane, at four different β values – $\beta = 2.5, 6.0, 7.0, 8.5$ (from top to bottom) – are shown to the right. The attractors shown on the u_f - \dot{q} plane, are also color-coded according to their type. On the top-left inset figure is shown a zoomed-in plot of $\langle J_{ac} \rangle$ -vs.- β at $\beta = 7.0$	66
2-3	The first 20 LEs at $\beta = 7.0$ and $\tau = 0.2$. The QR factorization is performed every timestep, ie, at a segment length of 0.01. The total integration time considered is 200 time units. Inset: the first 3 LEs at $\beta = 7.0$ and $\tau = 0.2$. The values obtained are $\lambda_1 \approx 0.19$, $\lambda_2 \approx 0.05$, and $\lambda_3 \approx -0.07$	68
2-4	Ergodic average of the Rayleigh index, $\langle J \rangle_{ray}$. Inset: zoomed-in plot of $\langle J \rangle_{ray}$ -vs.- β at $\beta = 7.0$ in the chaotic regime.	70

2-5	Ergodic average of the angles between the first 6 different adjoint CLVs (right) and tangent CLVs (left). The averaging window was set at 250 time units. The colorbar shows the angle in degrees.	72
2-6	Angles between pairs of adjoint and tangent CLVs when averaged over 250 time units. As expected, each tangent CLV is orthogonal to every adjoint CLV except those with the same LE.	73
2-7	Sensitivities of $\langle J_{ac} \rangle$ (red) and $\langle J_{ray} \rangle$ (blue) computed through tangent shadowing. The sensitivities are obtained by a cumulative average over sensitivities each of which is computed over 20 time units.	74
2-8	Sensitivities of $\langle J_{ac} \rangle$ with respect to τ (green, right) and $\langle J_{ac} \rangle$ (red, left), wrt β , computed through adjoint shadowing at $\beta = 6.9, \tau = 0.2$. The recorded sensitivities were obtained by a cumulative average over sensitivities each computed over a time length of 20 units.	75
2-9	The path of optimization starting at $\beta = 6.5$. The points are superimposed on the plot of $\langle J_{ac} \rangle$. In the gradient descent algorithm, the step size is taken to be 0.1.	76
2-10	Relative error in the instantaneous acoustic energy between its predicted and observed values in the chaotic regime ($\beta = 7$). The maximum error increases from blue to green colors. The mean error (over the assimilation time window) is shown in black.	82
2-11	Sensitivities $d\langle z \rangle / ds$ computed for the Lorenz'63 model using the tangent algorithm (left) and the adjoint algorithm (right). Different trajectories of length 3000 or 15 timeunits are used to compute the sensitivities, shown as an errorbar of length one standard deviation.	86
2-12	Relative error in the predicted state z as a function of time for the Lorenz'63 system. The maximum prediction error increases from light to dark colors. The mean error across all experiments is shown in black.	86

3-1	L: An ensemble of initial conditions distributed uniformly in the box $u_1 \in [0, 1], u_2 \in [0, 1], u_3 \in [28, 39]$ after 10 time units of evolution, shown on the $u_1 - u_3$ plane. R: distribution of a trajectory of 1000 time units in length. In both plots, the color represents the number of samples in a 2048×2048 uniform grid. The trajectory is sampled every 0.001 time units.	96
3-2	L: distribution of the same ensemble as in Figure 3-1 after another 5 time units of evolution. R: distribution of the same trajectory as in Figure 3-1 evolved for 10,000 time units.	96
3-3	L: distribution of the same ensemble as in Figures 3-1 and 3-2 after a total of 50 time units of evolution. R: distribution of the same trajectory as in Figures 3-1 and 3-2 evolved for 100,000 time units.	97
3-4	Periodic solutions of the Lorenz equation, overlaid on top of its SRB distribution. The solid line, dashed line, and dotted line represent three distinct periodic solutions.	99
3-5	L: the Lorenz map. The x-axis is the n -th local maximum of $u_3(t)$ over a long solution; the y-axis is the $(n + 1)$ -th local maximum of $u_3(t)$. The intersection of this curve with the dashed line ($y = x$) indicates the initial condition for the solid line in Figure 3-4. R: the Lorenz map iterated twice. The x-axis is the n -th local maximum of $u_3(t)$; the y-axis is the $(n + 2)$ -th maximum. The intersections with the diagonal dashed line indicate the initial conditions for both the dotted and dashed lines in Figure 3-4.	101
3-6	The tent map φ	102
3-7	Empirical probability distribution of long solutions (a billion steps) starting from four points whose binary digits have probability p of repeating the previous digit. The solution shown in the top-left plot starts from a point with $p = 0.5001$; top-right: $p = 0.51$; bottom-left: $p = 0.55$; bottom-right: $p = 0.9$	105

3-8	Empirical distribution function of other quasi-physical solutions. The solution shown in the left plot starts from an initial condition whose bits are independent and have probability 0.6 of being 0. The solution shown in the right plot starts from an initial condition whose bits have a probability of 0.6 of being 1 only following two consecutive 1's; otherwise a bit is 0 or 1 with equal probability.	107
3-9	L: The scaled tent map $\hat{\varphi}_s$ at different values of s between 0 and 1. R: sensitivity to small perturbation in the governing equation. The y-axis shows the absolute value of the difference between two solutions, one satisfying Eq. 3.12, one for $s = 0$ and the other for $s = 10^{-5}$. The x-axis shows the iteration number. The initial condition is at $x_0 = \pi/2$.	108
3-10	L: the tilted tent map Eq. 3.15 for $s = 0, 0.1, 0.2, 0.3, 0.4$, and 0.5 . R: the conjugacy \tilde{h}_s between the tent map (Eq. 3.6) and the tilted tent map (Eq. 3.15), evaluated using Eq. 3.16-3.17, for the same set of s as the left plot.	111
3-11	L: the density of a trajectory satisfying Eq. 3.6 that shadows a random trajectory satisfying Eq. 3.15 for $s = 0.1$. R: the density of a trajectory satisfying Eq. 3.6 that shadows a random trajectory satisfying Eq. 3.15 for $s = 0.5$	112
3-12	The “squashed” tent map (Eq. 3.18) for $s = 0.01, 0.05, 0.2$, and 0.5 . .	114
3-13	The left column shows the empirical probability distribution of the shadowing solution for the squashed tent map at $s = 0.01$ (top) and at $s = 0.2$ (bottom). The right column shows the empirical probability distribution of a physical solution at the same two values of $s = 0.01$ (top) and $s = 0.2$ (bottom).	115
3-14	The pinched tent map (Eq. 3.19) for $s = 0.01$ (blue), $s = 0.05$ (orange), $s = 0.2$, (green) and $s = 0.5$ (red).	116

3-15	The left column shows the empirical distribution of the shadowing solution for the pinched tent map at $s = 0.05$ (top) and at $s = 0.2$ (bottom). The right column shows the probability distribution of a physical solution at the same two values of $s = 0.05$ (top) and $s = 0.2$ (bottom).	118
3-16	The wave tent map (Eq. 3.20) for $s = 0.01, 0.05, 0.2,$ and 0.5	119
3-17	The left column shows the empirical distribution of the shadowing solution for the wave tent map (Eq. 3.20) at $s = 0.05$ (top) and at $s = 0.2$ (bottom). The right column shows the probability distribution of a physical solution at the same two values of $s = 0.05$ (top) and $s = 0.2$ (bottom).	120
3-18	The left, center, and right plots show the effect of the parameters $\sigma, \rho,$ and β on the Lorenz map, respectively. In the left plot, the blue and orange lines represent $\sigma = 10$ and $12,$ respectively, while $\rho = 28$ and $\beta = 8/3$. In the center plot, the blue and orange lines represent $\rho = 28$ and $30,$ respectively, while $\sigma = 10$ and $\beta = 8/3$. In the right plot, the blue and orange lines represent $\sigma = 8/3$ and $10/3$ respectively, while $\sigma = 10$ and $\rho = 28$	121
3-19	The left column shows the empirical distribution of the shadowing solution, and the right column shows the physical distribution of the Lorenz map at the following sets of parameters: top row: $\sigma = 15, \rho = 28, \beta = 8/3, \rho = 29,$ middle row: $\sigma = 10, \rho = 30, \beta = 10/3,$ and bottom row: $\sigma = 10, \rho = 28, \beta = 10/3$	122
3-20	The effect of increasing n on the physical probability distribution associated to the plucked tent map. The left column shows the plucked tent map at different values of s and $n = 0$ (top), $n = 3$ (middle), $n = 6$ (bottom). The original tent map is at $s = 0$ on each plot. The right column shows the stationary, physical probability distribution of the plucked tent map at $s = 0.1$ and $n = 0$ (top), $n = 3$ (middle), $n = 6$ (bottom).	125

3-21	L: the blue and orange lines are the Lorenz map at the standard parameters, and the dotted black lines indicate the approximate Lorenz map (Table 3.2). R: the regression error is shown as a function of z .	127
4-1	The mean squared error as a function of integration time τ , for different values of the variables k and r . The optimal value τ^* is marked on each of the plots.	146
4-2	Left: \tilde{e}_{\min}/k as a function of T , at different values of r . Right: Rate of convergence as a function of the ratio $r = \gamma_1/\lambda_1$	148
4-3	Left: Estimates of the variance and the bias of $\theta_{\tau,N}$ as a function of τ , for the Lorenz '63 system outlined in section 4.4.1. The dashed line indicates the least-squares fit over $\tau \leq 1.5$. Right: Sample mean estimates of $\mu_s(\theta_{\tau,N}^2)$ as a function of τ for the Lorenz '63 system outlined in section 4.4.1. The dashed line indicates the least-squares fit over $\tau \leq 1.5$	152
4-4	The variance and the square of the bias in the ensemble tangent sensitivity estimator for the Lorenz '96 system. The dashed line indicates the least-squares fit of the variance data.	154
4-5	Tangent density field v_ρ at time $t = 78 c/a_\infty$	156
4-6	Time evolution of L^2 norm of the tangent fields corresponding to the density ρ (top) and x -momentum ρu (bottom).	157
4-7	Left: Adjoint field corresponding to the density, labelled rhoa in the colormap, at $t^* = 0.35$. Right: The L^2 norm of the adjoint vector field, y , for the Navier-Stokes system in section 4.4.4 as a function of time.	159
5-1	Comparison of the sensitivities computed with pre-S3 to finite-difference for the solenoid map in Section 5.7.1. (a) J is a set of two-variable nodal basis functions along r and θ axes. (b) J_θ is a set of nodal basis functions along θ axis.	190

5-2	Comparison of the sensitivities of the nodal basis functions along the θ and ϕ axes to the parameter s_2 obtained for the Kuznetsov-Plykin attractor using (a) finite difference and (b) the pre-S3 algorithm. . . .	191
6-1	Comparison of the x_1, x_2, x_3 components of V^1 computed analytically (orange circles) and numerically (blue crosses), for the super-contracting Solenoid map	215
6-2	Comparison of the x_1, x_2, x_3 components of W^1 computed analytically (orange circles) and numerically (blue crosses), for the super-contracting Solenoid map.	215
6-3	The vector field V^1 is shown for the Solenoid map. The color represents $\ W^1\ $, which is the curvature of the attractor.	216
6-4	Orbit points of the Lorenz system shown on the x_1 - x_3 plane, at $T_1 = 18$ (left) and at $T_2 = 20$ (right), colored according to the distance from their centroid normalized by the centroid z -coordinate. The initial conditions were 10001 equi-spaced points on the short line segment joining $(-0.01,0,1)$ and $(0.01,0,1)$	217
6-5	Comparison between V^1 from an iteration of the tangent dynamics (shown in orange) and V^1 from finite difference of the primal trajectories (in blue). The first column shows the components of V^1 at time $T_1 = 18$ and the second column at $T_2 = 20$. The first, second and third rows show the x_1, x_2, x_3 components of V^1 respectively.	218
6-6	Comparison between W^1 from the differential CLV method (shown in orange) and W^1 from finite difference (in gray). The first column shows the components of W^1 at time $T_1 = 18$ and the second column at $T_2 = 20$. The first, second and third rows show the x_1, x_2, x_3 components of W^1 respectively.	219
6-7	The vector fields V^1 (left) and V^2 (right) are shown for the PCM at $s_1 = 0.75, s_2 = 0.2$	223

6-8	The vector field V^1 is shown for the PCM at $s_1 = 0.75, s_2 = 0.2$. The color represents the values of $\ W^1 \times V^1\ $, which equals the norm of $\ W^1\ $ multiplied by a sign representing the orientation with respect to V^1	223
6-9	The vector field V^1 is shown for the DCM at different parameter choices. The parameters not indicated are set to 0 in each case. . . .	225
6-10	The vector field V^1 is shown for the DCM, colored according to $\ W^1 \times V^1\ $. The parameters not indicated as 1 are set to zero in each case.	226
6-11	The CLV V^1 on the henon attractor. Inset is the CLV field in a neighborhood of the fixed point $\approx (0.63, 0.19)$	228
6-12	The vector field V^1 is shown for the Hènon map. The color represents the V^1 self-derivative norm, $\ W^1\ $	229
6-13	The vector field V^1 is shown for the Hènon map. The color represents $\ W^1\ $, the curvature of the unstable manifold.	229
7-1	Top left: the domain \mathbb{T}^2 covered by rectangles. The other figures show the application of Baker's map at $s = 0 \in \mathbb{R}^4$ (top right), $s = [0.2, 0, 0, 0]^T$ (bottom left) and $s = [0, 0.2, 0, 0]$ (bottom right) on the gridded top left figure.	273
7-2	Top left: the domain \mathbb{T}^2 covered by rectangles. The other figures show the application of Baker's map at $s = 0 \in \mathbb{R}^4$ (top right), $s = [0, 0, 0.2, 0]^T$ (bottom left) and $s = [0, 0, 0, 0.2]$ (bottom right) on the gridded top left figure.	274
7-3	Each plot shows the SRB distribution achieved at the parameter values indicated on the title.	275
7-4	Ergodic average of the objective function $J = \cos(4x_2)$ as a function of s_4 (left) and s_1 (right); the other parameters are set to 0. Sensitivities from S3 are shown in black at select parameter values.	276

7-5	Ergodic average of the objective function $J = \cos(4x_2)$ as a function of s_4 (left) and $s_1 = s_3$ (right); the other parameters are set to 0. Sensitivities from S3 are shown in black at select parameter values.	276
8-1	The parametric derivatives $d_{s_2}\langle z \rangle$ (left) of the Lorenz system with error bars indicating one standard deviation. On the right, $d_{s_2}\langle (z - 28)^2 \rangle$ computed using S3 at various different s_2 values on the x -axis are shown as black lines. The blue points on the right hand side subfigure are computed as ergodic averages along trajectories of length 1.6 billion (3.2 million time units). The other two parameters s_1 and s_3 are held at their standard values.	312
9-1	The projection of the Solenoid attractor on the $x^{(1)}-x^{(2)}$ plane at $s = [1, 0]^T$ (left), $s = [1, 1]^T$ (center) and $s = [2, 0]^T$ (right) respectively.	316
9-2	The sensitivities of $\langle J \rangle = \langle x^{(1)}.x^{(1)} + x^{(2)}.x^{(2)} \rangle$ with respect to the parameter s_1 (left) and s_2 (right) computed by S3 are shown as black lines at a number of parameter values on the $\langle J \rangle$ vs parameter curve.	317
9-3	Comparison of relative error in the sensitivity of $\langle J \rangle$ to s_2 at $s = [1, 0]$ of the Solenoid map (section 9.1) as a function of the number of samples, between S3 and finite difference. The relative error is computed with respect to the true value, which is taken to be the S3 derivative obtained with $N = 9.6$ billion samples. The finite difference in the parameter s_2 used to calculate the derivatives is set to 0.001 in order to obtain the red dots. The green triangles are the derivatives computed using the S3 algorithm described in section 7.3.1. The dotted lines have a slope of -0.5.	320

List of Tables

3.1	Comparison of statistics computed from periodic solutions of the Lorenz equation with the statistics computed from physical solutions	100
3.2	The fitting parameters of the Lorenz map at $\sigma = 10, \beta = 8/3$, and $\rho = 28$	128
9.1	Summary of the cost of the S3 algorithm as described in section 7.3.1 and its reverse-mode/adjoint version.	325

Chapter 1

Introduction

“Oh, and you know the thing about chaos? It’s fair!”, The Joker, in *The Dark Knight* (taken entirely out of context).

This thesis is a step toward the solution of an important computational problem concerning chaotic dynamical systems: an efficient estimation of their long-term response to small perturbations. In this introductory chapter, we explain this problem in detail and describe the cross-cutting applications of a practical algorithmic solution to this problem. We examine the computational and mathematical difficulties posed by chaos, which have hampered the application of existing methods for sensitivity analysis of dynamical systems.

1.1 Long-term or statistical response to infinitesimal perturbations and its applications

A common question of both fundamental scientific and practical importance that arises in virtually every field of science and engineering is the following: how does the long-term behavior of a dynamical system respond to small perturbations of

system parameters? This question falls within the purview of sensitivity analysis, which entails the computation of derivatives of outputs of interest obtained from a dynamical system with respect to user-defined inputs, such as system parameters. Such derivatives or sensitivities are obtained by solving linear perturbation equations, i.e., tangent or adjoint equations [94][95] numerically or by automatic differentiation [34][227], or by using finite-difference methods [194]. The resulting sensitivities are routinely used, when the dynamical system is non-chaotic, in gradient-based computational approaches for various target applications such as optimization and design [111][112][201][31][140][203], control, analysis and prediction [196][229][102][103][82][215], uncertainty quantification [104][270][216][7][192][283], data assimilation [95][269][131][195][24][125], parameter selection [52][8][53][34][33] and inverse problems [151][255][254][115][258][151]. The vast literature on sensitivity analysis applications that we have only sparsely sampled here indicates not only that the computed sensitivities are indispensable for various engineering tasks performed by such gradient-based target computations, but that they are truly cross-cutting quantities of interest. Even in the small collection we refer to above, gradient-based computational approaches are used crucially in diverse disciplines such as biomedical engineering, climate science and aerospace engineering.

Such full-fledged applications of sensitivity computation have so far been restricted to situations wherein the underlying dynamical model is non-chaotic. When the dynamical model exhibits chaos, and when the sensitivities needed are of long-time averaged or equivalently, statistical quantities, such conventional linear perturbation methods fail. This failure is a manifestation of the so-called *butterfly effect*, which is the defining feature of chaotic dynamics: an infinitesimal perturbation to the state grows exponentially in norm with time. For example, consider the solutions of linear perturbation equations, computed using tangent, adjoint, finite-difference methods shown in Figure 1-1. The underlying dynamical system, which is treated in Chapter 2, models chaotic thermoacoustic instabilities in a Rijke tube/acoustic cavity, which serves as a reduced-order model for a gas turbine engine.

The plots indicate that an infinitesimal parameter perturbation, which is an in-




Figure 1-1: Linear perturbation solutions computed for the thermoacoustic system discussed in Chapter 2. FD stands for finite-difference and AD for automatic differentiation.

infinitesimal perturbation to the state at each time, causes an exponential divergence from the unperturbed trajectory (at the original value of the parameter). This exponential divergence also holds when automatic differentiation is used to compute the tangent or adjoint solutions, as shown in Figure 1-1.

On the other hand, in many physical systems that obey the chaotic hypothesis [236][108], the statistics or long-time averages of the systems vary differentiably with respect to parameters. In other words, the derivative of the long-time averages or ensemble averages of these systems with respect to system parameters are bounded quantities, although along almost every trajectory, the derivative of the time-average increases exponentially with time and does not converge. If s is a system parameter that is varied by an infinitesimal amount δs , and $\langle J \rangle(s)$ is the ensemble average of a state function J , $\langle J \rangle(s + \delta s) = \langle J \rangle(s) + (\delta s)d_s \langle J \rangle(s) + O(\delta s^2)$. The derivative $d_s \langle J \rangle(s)$ is referred to as *linear response* because this statistical response is linear in the perturbation δs . The existence of this derivative implies that, up to first order in the parameter perturbation, the ensemble/infinately-long time average of an observable can be expressed purely using information from the original, unperturbed dynamical

system.

A notion of ergodicity is implicitly assumed in the above discussion of linear response. Ergodicity of a dynamical system is a property that implies that its long-time behavior is independent of its initial condition. Starting from almost every initial state, the infinitely-long time average of an observable converges to the same value. This is why, although the observable J depends on the system state, $\langle J \rangle$ is only a function of the system parameters and independent of the initial condition of the dynamics. This common value of an infinitely-long time average is the ensemble average of the observable with respect to a particular stationary probability measure over states. In uniformly hyperbolic systems [46], which we treat exclusively in this thesis, such a stationary measure, known as Sinai-Ruelle-Bowen or *SRB* measure exists that captures the long-time behavior of almost every trajectory of the system. Thus, $\langle J \rangle$ is both the infinitely-long time average of the observable J along a single trajectory starting almost everywhere and the expectation of the observable of J with the states distributed according to the SRB measure. The crux of this thesis is a computational method to compute linear response, $d_s \langle J \rangle$, which is the parametric derivative of expectations according to the SRB measure. In uniformly hyperbolic systems, Ruelle [232][234][235] first showed that linear response holds, i.e., the statistics are differentiable with respect to parameters, and proved a rigorous formula for linear response. However, in its original form, the computation of the Ruelle’s response formula is inefficient [69][96], as we shall discuss in detail in Chapter 4.

An efficient computation of linear response will usher the entry of downstream computational applications mentioned before, such as design, optimization, uncertainty quantification, parameter selection, data assimilation, control theoretic and inverse problems, into the realm of chaotic dynamics. These downstream applications have been largely absent on chaotic models such as mechanical and aerospace engineering turbulent fluid flows and molecular simulations for biomedical applications, wherein long-time averaged/statistical behavior is of interest. For example, flow control, output error-based mesh adaptation for chaotic unsteady simulations, and design optimization for jet noise reduction, building engine components that are

robust to i) fatigue induced by chaotic aerodynamics, and ii) chaotic thermoacoustic instabilities, remain significant engineering challenges due to the unavailability of sensitivities [104][79][264].

Beyond having immense practical value, the computation of linear response, which gives us a quantitative measure of the influence of each parameter on the long-term behavior of a chaotic system, improves our fundamental understanding of chaotic phenomena, such as fluid turbulence. An important area where the computation of linear response in a chaotic model arises is climate sensitivity [228][185][9], wherein the response of the climate system – a chaotic dynamical system – to greenhouse gas concentrations, and the reduction in the uncertainty in this response (which in turn involves linear response calculations in chaotic models of microphysical processes) are actively sought [231][156][242][282].

1.2 Existing methods for linear response computation

A few methods exist that circumvent exponentially growing sensitivities to compute a finite linear response value. Ruelle’s linear response formula [232][234][235] expresses the desired derivative as a series summation of ensemble averages of short-term sensitivities. Ensemble sensitivity methods, due to Lea *et al.* [168], are a direct approximation of Ruelle’s formula wherein each term of the series is approximated as a sample average of short-time sensitivities computed using previously mentioned linear perturbation methods. However, the number of samples needed to reduce the variance in the exponentially growing sensitivities and compute linear response accurately makes this approach computationally infeasible [96]. In blended response algorithms [1][2], the ensemble sensitivity approach for short-time sensitivities is *blended* with a fluctuation-dissipation theorem-based approximation of the long-term sensitivities. This approximation is however adhoc since the densities of the SRB measure on unstable manifolds may not follow the fluctuation-dissipation theorem-based approx-

imation, while linear response still holds.

A recent work by Ni [210] introduces the linear response algorithm, which appears to be another viable solution to linear response computation. It has been derived for uniformly hyperbolic systems with arbitrary dimensional unstable manifolds. This algorithm provides a “fast” computation of the unstable divergence via a recursive formula, analogous to the density gradient for one-dimensional unstable manifolds treated in Chapter 7. In this linear response algorithm, second-order tangent equations, which are the most expensive step, are solved to differentiate certain vector fields with respect to a modified shadowing direction. Unlike in Ni [210], the second-order tangent equations developed in the present thesis (Chapter 6 and Chapter 7) are derivatives along the one-dimensional unstable manifold.

The shadowing direction is computed by the non-intrusive least squares shadowing algorithm [206][211][271], which was independently proposed as a method to compute linear response. Least-squares shadowing [271][72][273] and its non-intrusive variants [206][211][37] compute linear perturbation solutions – *shadowing* perturbations – that remain bounded for a long time. While these lead to bounded sensitivities, these are not unbiased estimates of linear response [209]. We explore the probability distribution of a shadowing solution, which is the foundation of these shadowing sensitivity analysis methods, in Chapter 3.

1.3 Thesis contributions

As the main contribution of this work, we introduce the space-split sensitivity or S3 algorithm to compute linear response, or the parametric derivative of statistics, in chaotic systems. In order to be applicable to high-dimensional chaotic numerical simulations, the S3 algorithm is designed to compute linear response as a time average along a long trajectory of the system, which is the solution of a chaotic ODE/PDE over a long time. The S3 algorithm provably converges, as the trajectory length tends to infinity, to Ruelle’s linear response formula. S3 is an exact computation of linear response that has an error convergence of $\mathcal{O}(\sqrt{\log \log N}/\sqrt{N})$, almost like a typical

Monte Carlo integration, when computed with an N -length simulation, independent of the system dimension.

A secondary purpose of this thesis is the analysis of some of the previous attempts at linear response computation, including the shadowing algorithm [271][206] and ensemble sensitivity computation [168]. We perform a detailed feasibility analysis of the ensemble sensitivity method in Chapter 4 (also published as [69]), and conclude, in agreement with previous works [96], that this method is, in general, impractical in high-dimensional chaotic systems. In Chapter 3, we give examples to prove that shadowing trajectories can be non-physical. This analysis reveals that the shadowing algorithm is not guaranteed to converge to the true linear response value. The analysis also dispels the notion that small parameter perturbations can result only in small changes in statistics of chaotic systems [61].

1.4 Outline

Each remaining chapter is self-contained, and the notation and definitions are not shared except among Chapters 7, 8 and 10. The organization of the chapters is as follows. While shadowing is not guaranteed to give the correct value of sensitivities in general, Chapter 2 treats an example where shadowing sensitivities are accurate enough to yield meaningful downstream applications. We consider a reduced-order model of an acoustic cavity that exhibits chaotic thermoacoustic instabilities. We perform optimization and introduce a new data assimilation scheme both of which use sensitivities computed by the shadowing algorithm. In Chapter 3, we prove through examples that shadowing orbits can be nonphysical and discuss the implications of this result for numerical simulation and sensitivity analysis of chaotic dynamical systems. Turning our attention to Ruelle's formula, in Chapter 4, we provide a feasibility analysis of its direction evaluation – the ensemble sensitivity method.

In Chapter 5, a first version of the S3 (space-split sensitivity) algorithm is presented in which Ruelle's formula is decomposed by splitting the parameter perturbation into its stable and unstable components. For an expanding map in which the

parameter perturbation would be purely unstable, Chapter 5 provides a derivation of an alternative formulation of linear response that is efficiently computable. A recursion trick is used to modify Ruelle’s formula to a form amenable to ergodic-averaging. This chapter assumes the differentiability of the stable and unstable vector fields on phase space, which is in general false [130], and the final version of S3, in Chapter 7, does not make this assumption.

Chapter 7 develops the S3 algorithm for a generic uniformly hyperbolic system with a one-dimensional unstable manifold. The basic idea is a particular decomposition of Ruelle’s formula that involves a regularized tangent equation solution. The regularized tangent solution leads to a decomposition into two components, which, although still named stable and unstable contributions, differ from that in Chapter 5. We prove that such a solution that leads to a computable reformulation of Ruelle’s formula exists and is differentiable on the unstable manifold. This chapter contains the proof of this decomposition, known as the S3 decomposition, and the proof of convergence of the S3 algorithm. The resulting unstable contribution involves the derivative of the unstable direction along itself. A computational algorithm for such derivatives is the subject of Chapter 6.

Chapter 8 extends the S3 algorithm presented for uniformly hyperbolic maps in the previous chapter to uniformly hyperbolic flows that additionally have a center direction. This situation is exemplified by time-discretized ODEs in which case the center direction is spanned by the generating vector field. The S3 algorithm now includes a center contribution, in addition to the stable and unstable contributions, which are modified slightly from the case of uniformly hyperbolic maps. A detailed cost and error analysis of the S3 algorithm is carried out in Chapter 9. The final chapter presents a brief road map for future work on extending the S3 algorithm to maps with arbitrary dimensional unstable manifolds. The algorithm developed in this thesis for one-dimensional unstable manifolds serves as a starting point.

Chapter 2

Shadowing and its application to a time-delayed acoustic model

In this chapter, we perform sensitivity analysis of long-time (or ensemble) averages in the chaotic regime using the shadowing algorithm. We introduce automatic differentiation to eliminate the tangent/adjoint equation solvers used in the shadowing algorithm. In a gradient-based optimization, we use the computed shadowing sensitivity to minimize different long-time averaged functionals of a chaotic time delayed system by optimal parameter selection. In combined state and parameter estimation for data assimilation, we use the computed sensitivity to predict the optimal trajectory given information from a model and data from measurements beyond the predictability time. The algorithms are applied to a thermoacoustic model. Because the computational framework is rather general, the techniques presented in this chapter may be used for sensitivity analysis of ensemble averages, parameter optimization and data assimilation of other chaotic problems, where shadowing methods are applicable.

This chapter also serves a tutorial on the shadowing algorithm [271][211][37], the introduction of automatic differentiation into the algorithm and its application to optimization and data assimilation. It is a joint work with Luca Magri and Qiqi Wang and is available as a preprint at [58].

2.1 Introduction

Sensitivities are quantitative measures of the response of model outputs to infinitesimal changes in inputs, which are crucial to engineering design [251, 161]. They are employed in parameter estimation and model selection [11, 147], uncertainty quantification [4, 177], data assimilation [239, 164, 256] and design and optimization (see [237, 45, 268, 190] for recent reviews of the applications of sensitivity derivatives in different engineering disciplines). With the growing ability to simulate high-dimensional complex dynamics, much research effort has been invested into commensurately improving sensitivity analysis methods. Adjoint of mathematical models have been developed and used successfully in many fields; for example, adjoint sensitivity analysis in meteorology [110, 134], aircraft design [250], systems biology [155], chemical kinetics [239], thermo-fluids [190], among others. The adjoint method is generally used when the input parameter space is high-dimensional. Tangent Linear Models (TLMs) and finite difference methods are also used for sensitivity analysis when the dimension of the input parameter space is small enough that the cost of simulating the original dynamics repeatedly is not prohibitive. In many of these applications, automatic differentiation (AD) has successfully replaced the TLM or adjoint computations [48, 121]. For example, MITGCM [3], a popular climate model uses `OpenAD` [278, 133], an open source source-transformation AD software [266, 265, 6, 138]; and `Tapenade` [128, 129] has replaced adjoint differentiation in a few industrial-size numerical codes.

In recent times, simultaneous advances in simulation capabilities and computing power have led to a proliferation of scale-resolving simulations of chaotic systems [39, 47, 92, 252]. For many of the above target applications, the relevant observables, or outputs, in chaotic systems are statistically stationary or infinitely long-time averaged functions [47, 92, 252, 141, 142]. Useful gradients of ensemble averages, which are equal to infinite time averages in ergodic systems, cannot be obtained by time-averaging the instantaneous gradients in chaotic systems [271]. Indeed, in the infinite time limit, the time averages of the instantaneous gradients diverge despite that the

ensemble averages of the functions may have a well-defined gradient. This is because the tangent space of a chaotic attractor is exponentially unstable. Likewise, the sensitivities computed on time-integrating the adjoint sensitivities diverge exponentially. For the same reasons, other methods, such as TLMs and AD, also fail to compute meaningful sensitivities in chaotic systems. Due to these challenges, sensitivity analysis of chaotic systems has not developed as much as sensitivity analysis of non-chaotic systems has [206, 208, 271].

One approach for the computation of sensitivities of long-time-averaged functionals in chaotic systems is the Least Squares Shadowing [271] (LSS) method. This method by-passes the exponential instability by computing the derivatives along a close *shadowing* direction, which is obtained as the solution of a constrained minimization problem. A recent variant of the LSS method—the Non-Intrusive Least Squares Shadowing (NILSS) [206]—has been proposed to reduce the computational cost and memory requirements of the original LSS problem by projecting the gradients onto the *unstable* subspaces only. The NILSS method has found successful applications in chaotic computational fluid dynamics; e.g., [39] and [208] applied NILSS to scale-resolving Direct Simulations of chaotic flows around bluff bodies, [37] developed the adjoint version of NILSS and applied it to wall-bounded chaotic flows; and [141] applied it to the optimization of chaotic acoustic oscillations subject to synthetic turbulence. Other methods for sensitivity analysis in chaotic systems are conceptually based on extensions of the fluctuation-dissipation theorem for nonequilibrium systems in physics. One method is based on estimating the invariant probability distribution [80]. Other recent approaches [2, 77, 185, 2] computationally evaluate Ruelle’s response formula for nonequilibrium systems [233]. In this chapter, we are concerned with the NILSS algorithm and, in particular, on the development of the automatically differentiated version of the algorithm. The algorithm is generalized to tackle time delayed systems.

Delayed differential systems, which often tend to be chaotic, are extensively used for mathematical modelling of transport and non-Markovian processes, such as population dynamics and cell proliferation in mathematical ecology and biology [158, 249],

chemical processes [51], neural networks, networked control systems [284, 30] thermoacoustics [190], among others. Sensitivity derivatives with respect to the parameters, including the delay parameter, have been employed for model selection, system identification and stability analysis [23, 143, 190]. The objective function to be optimized, in the case of the model parameter estimation or system identification problems, is typically a mean quantity that depends on the parameters, including the time delays [23, 141]. The tangent and adjoint discrete AD-shadowing methods developed in this chapter offer a solution to this problem in chaotic time-delayed systems. In order to compute sensitivities using AD, differentiating the numerical solution of the primal dynamics is necessary. If the time delay is treated approximately as an integer number of timesteps by the time integrator, differentiating with respect to the delay poses a problem since chain rule differentiability is lost. In order to circumvent this issue, in this chapter we exploit the fact that the effect of a time delay can be replaced by a linear advection equation, and, therefore, we solve for an extended primal system.

A practical engineering problem modelled with time-delayed equations is thermoacoustics [190]. Gas-turbine and rocket-motor manufacturers strive to design engines that do not experience thermoacoustic instabilities [171]. Thermoacoustic instabilities occur when the heat released by the flame is sufficiently in phase with the acoustic pressure [178], such that the thermal energy of the flame that is converted into acoustic energy exceeds dissipation mechanisms. Unstable thermoacoustic systems have intricate nonlinear behaviours when design parameters are varied, from periodic, through quasi periodic to chaotic oscillations [149]. Although methods to investigate the sensitivity of fixed points (with eigenvalue analysis) and periodic solutions (with Floquet analysis) are well-established [190], a stability and sensitivity framework to tackle chaotic acoustic oscillations is only at its infancy [141, 142]. In thermoacoustics, sensitivity analysis quantitatively informs the practitioner on how to optimally change design parameters, such as geometric quantities; which passive device is most stabilizing; and how large is the uncertainty of the stability calculations [189, 243, 199], as reviewed by [190]. All these studies are concerned with the calculation of sensitivities of eigenvalues around non-chaotic attractors. These established eigenvalue-sensitivity

methods fail in chaotic systems because of the butterfly effect [69, 168, 96, 141] (§2.2). In this chapter, we apply the computational framework we develop to the calculation of the derivative of two infinite time-averaged cost functionals, one being an energy norm and the second being an integral metric, with respect to the parameters' vector. These derivatives give us a quantitative estimate of the long-term response of chaotic acoustic oscillations. We use these sensitivities to stabilize a nonlinearly unstable, yet eigenvalue-stable, thermoacoustic system. Physically, the cost functionals represent the acoustic energy, which we want to minimize to make the combustor operate in stable conditions. We use these sensitivities in a gradient-based optimization algorithm to suppress a chaotic acoustic oscillation, which cannot be achieved by only stabilizing the eigenvalues or through short-term chaotic sensitivity calculations.

The chapter is structured as follows. In section 2.2.5, the idea behind the NILSS algorithm is reviewed. Section 2.2 defines the problem with a mathematical background on sensitivity analysis in chaotic systems. The main features of the shadowing algorithm are explained in 2.3. The AD version of the algorithm is provided in 2.10. We present the chaotic time-delayed model of a prototypical thermoacoustic system in section 2.4. The tangent and adjoint shadowing sensitivities of this model are calculated and applied for parameter estimation for gradient-based optimization in section 2.5, and for data assimilation in section 2.6. The chapter ends with a final discussion in section 2.7.

2.2 Shadowing sensitivity in chaotic systems

Before we describe the NILSS algorithm, we recall the problem of extreme sensitivity to perturbations in chaotic systems, which leads to ill-conditioning of linearized models, such as the tangent equation, the adjoint equation and algorithmic differentiation. We define the primal problem by a set of ordinary differential equations

(ODEs), which may be spatially discretized partial differential equations, as,

$$\begin{aligned}\frac{du}{dt} &= \mathcal{F}(u, \mathcal{S}), \quad \mathcal{S} \in \mathbb{R}^p \\ u(0) &= u_0 \in \mathbb{R}^d.\end{aligned}\tag{2.1}$$

Here $u \in \mathbb{R}^d$ is the state of the system, and, $\mathcal{S} \in \mathbb{R}^p$ is a vector of system parameters. The system parameters, which can be, e.g., control variables in an adjoint-based design problem, do not change with time. The right hand side $\mathcal{F} : \mathbb{R}^d \times \mathbb{R}^p \rightarrow \mathbb{R}^d$ of the primal ODE (Eq. 2.1), which is also referred to as the *time-derivative* direction, is a function of the instantaneous state and \mathcal{S} . In this chapter, we study the discrete-time system obtained by time-integration of the primal ODE. Throughout, we use a subscript to denote a discrete time, which is represented by a positive integer. In particular, $u_0 \in \mathbb{R}^d$ is the initial state; $u_n \in \mathbb{R}^d$ is the solution vector at time $n \in \mathbb{Z}^+$.

We define the function $f : \mathbb{R}^d \times \mathbb{R}^p \rightarrow \mathbb{R}^d$ to denote the time-one map, i.e., the time-integrator that evolves a solution state by one timestep, so that $u_1 = f(u_0, \mathcal{S})$. We use the notation f_n to denote the n -time composition of the map f , at a fixed set of parameters, so that $u_n = f_n(u, \mathcal{S})$, $n \in \mathbb{Z}^+$. The set of vectors $\{u_n\}$ is an *orbit* or a *trajectory* of the dynamics $f(\cdot, \mathcal{S})$. Let \mathcal{J} be a set of l scalar observables in $\mathcal{C}^2(\mathbb{R}^d)$, and J be an observable in this set. Given an initial state u_0 , the N -time average of J is $\langle J \rangle_N := (1/N) \sum_{n=0}^{N-1} J(u_n)$. In ergodic systems, in the limit $N \rightarrow \infty$, the N -time-average, which is referred to as *ergodic average* and denoted as $\langle J \rangle$, is well-defined and independent of the initial state u_0 . The ergodic average $\langle J \rangle$ is a function of the parameters \mathcal{S} only. Its value is equal to an expectation of J with respect to the ergodic, stationary probability distribution achieved by the state vector under the dynamics f . In chaotic systems, ergodic averages of observables are often the quantities of interest for optimization and control problems. In these problems, the long-term response of a chaotic system to infinitesimal perturbations may be desired (e.g. [184]), as opposed to a short-term or intermediate-term response. The problem of nonlinear acoustic oscillations that is studied in this chapter is one such example

[141]. Our goal is to compute, for all $J \in \mathcal{J}$ and all $s \in \mathcal{S}$ the quantity,

$$d_s \langle J \rangle := d_s \left(\lim_{N \rightarrow \infty} \langle J \rangle_N \right), \quad (2.2)$$

where $d_s := d/ds$ denotes the differentiation operator with respect to s . We assume that the ergodic average $\langle J \rangle$ is differentiable with respect to s . In *uniformly hyperbolic* systems, a stationary probability distribution, known as the SRB measure [280], exists, with respect to which ergodic averages converge starting from an open set in \mathbb{R}^d containing the attractor. Under certain smoothness conditions on the map, the SRB measure is differentiable with respect to parameters [233], for small, smooth parameter perturbations. The assumption of uniform hyperbolicity is involved in the shadowing algorithm (section 2.3) and in the data assimilation scheme (section 2.6). There is a wealth of numerical and experimental evidence [107, 208] that shows that physical systems exhibit *quasi-hyperbolic* behavior. Hyperbolicity of the time-delayed system we consider in this chapter, has been numerically verified by Huhn and Magri [141] for a range of design parameters (numerical experiments are also presented in Figure 2-5 later in this chapter).

2.2.1 Tangent dynamics

The tangent equation describes the response of the system's state to infinitesimal perturbations in a parameter $s \in \mathcal{S}$ in a neighborhood of a reference trajectory $\{u_n\}$. By introducing the shorthand $v_n := \partial_s u_n$, the tangent equation is

$$\begin{aligned} v_{n+1} &= \partial_s f(u_n, \mathcal{S}) + (D_u f)(u_n, \mathcal{S}) v_n \\ v_0 &= 0 \in \mathbb{R}^d, \end{aligned} \quad (2.3)$$

where D_u denotes the differentiation with respect to the state vector. We refer to the solutions v_n as the inhomogeneous tangent solutions. On setting the *source* term in Eq. 2.3 to zero, and starting with a non-zero initial perturbation, we obtain the time evolution of the perturbations in the initial state, denoted q_n . We refer to q_n as the

homogeneous tangent solution whose time evolution is given by

$$q_{n+1} = (D_u f)(u_n, \mathcal{S}) q_n, \quad (2.4)$$

The solution q_n is the derivative: $q_n := (D_u f_n)(u_0, \mathcal{S})q_0$, which means that the homogeneous tangent equation is an iterative application of the chain rule. The homogeneous tangent solution is the difference between two orbits of f at n , which are separated by an infinitesimal distance along q_0 at time 0. The inhomogeneous tangent solution, on the other hand, is the difference between two orbits of $f(\cdot, \mathcal{S})$ at infinitesimally different s , starting from the same initial condition. We can write down the following difference approximation of the inhomogeneous tangent equation:

$$v_n \approx \frac{f_n(u, s + \epsilon) - u_n}{\epsilon}. \quad (2.5)$$

The homogeneous tangent equation can also be approximately computed by finite differences

$$q_n \approx \frac{f_n(u + \epsilon q_0, \mathcal{S}) - u_n}{\epsilon}, \quad (2.6)$$

Both these finite difference approximations are valid only up to an index n for which the perturbed trajectory, and the original trajectory $\{u_n\}$, remain near each other. Because of chaos, for almost every direction q_0 , the perturbed and unperturbed trajectories, exponentially diverge from each other. The finite difference approximations are bounded by D/ϵ , where D is a scalar upper bound for the attractor within which all state vectors lie. On the other hand, the tangent solutions v_n and q_n , which are limits as $\epsilon \rightarrow 0$ of the right hand sides of Eq. 2.5 and Eq. 2.6, respectively, continue to diverge exponentially with n in a chaotic system, unlike the finite difference approximations. That is, for large n , and almost every q_0 , $\|v_n\|, \|q_n\| \sim e^{\lambda_1 n}$, where $\lambda_1 > 0$ is the largest characteristic *Lyapunov exponent* [15]. In this chapter, we refer to $1/\lambda_1$ as the *Lyapunov time*, which is a timescale for the number of iterations needed to increase the norm of a linear perturbation by a factor e . However, not all infinitesimal

perturbations diverge exponentially. There are initial conditions, q_0 , at every u on a chaotic attractor that generate (asymptotically) exponentially decaying homogeneous tangent solutions along the orbit of u . This is because the space of tangent solutions \mathbb{R}^d , has the direct sum decomposition $\mathbb{R}^d = E^u(u) \oplus E^s(u) \oplus E^c(u)$. The linear subspaces $E^u(u)$ and $E^s(u)$ contain tangent vectors at u that exponentially grow/decay asymptotically under the tangent dynamics in Eq. 2.3, respectively. There is a maximum of d possible Lyapunov exponents that give the asymptotic exponential growth/decay rates of tangent vectors (see [15] for Oseledets theorem). Hereafter, we assume that there are d_u strictly positive Lyapunov exponents, which means that the unstable subspace at each point is d_u -dimensional. In a chaotic system, $d_u \geq 1$, i.e., the unstable subspace is at least one-dimensional at every u . In this chapter, the center subspace $E^c(u)$ consists of all the tangent vectors that asymptotically neither grow nor decay on an exponential scale, i.e. the tangent vectors in this subspace have a zero Lyapunov exponent. For example, consider the tangent vector $\mathcal{F}(u)$, whose flow is our primal system (Eq. 2.1). If f is a numerical discretization of the dynamical system in Eq. 2.1, \mathcal{F} approximately satisfies Eq. 2.3 (it exactly satisfies the continuous-in-time formulation of Eq. 2.3). In this chapter, we assume that $E^c(u)$ is one-dimensional at every u and is spanned by a bounded tangent vector field, say $\tilde{\mathcal{F}}$, which exactly satisfies Eq. 2.3,

$$\tilde{\mathcal{F}}(u_{n+1}) = (\partial_s f)(u_n, \mathcal{S}) + (D_u f)(u_n, \mathcal{S}) \tilde{\mathcal{F}}(u_n), \quad (2.7)$$

The vector field $\tilde{\mathcal{F}}$ is approximated by the known vector field \mathcal{F} , and will be referred to as the center direction. Note that this is a slight generalization of uniform hyperbolicity (E^c is technically absent in a uniformly hyperbolic system), for which we assume the uniqueness and differentiability of the SRB measure [221].

2.2.2 Adjoint dynamics

Exponential divergence also holds for adjoint equations starting from almost every initial condition. Fixing $N \in \mathbb{N}$, the N -time average $\langle J \rangle_N$ is affected by the primal

solution at each $n \leq N$. If u_n is infinitesimally perturbed, then, $J(u_m)$ is modified for all $m \geq n$, causing the N -time average $\langle J \rangle_N$ to be altered. Viewed in this manner, at a fixed \mathcal{S} , $\langle J \rangle_N$ is a function of N variables, $\{u_n\}_{n=0}^{N-1}$, where each variable u_m is, in turn, a function of u_n , $n < m$. That is, $\langle J \rangle_N \equiv \langle J \rangle_N(u_0, u_1, u_1, \dots, u_{N-1})$, with $u_n = f(u_{n-1}, \mathcal{S})$. The adjoint solution at time n is the response of $\langle J \rangle_N$ to an infinitesimal perturbation in u_n , keeping the states prior to n fixed at a reference orbit. At $n \leq N$, the adjoint solution is defined as

$$v_n^* := (D_{u_n} \langle J \rangle_N)^T(\{u_n\}_{n=0}^{N-1}) \in \mathbb{R}^d, \quad (2.8)$$

where D_{u_n} refers to the total derivative with respect to u_n and T stands for transpose. For $n = N$, Eq. 2.8 gives $v_N^* = (1/N) (D_u J)^T(u_N, \mathcal{S})$. By applying the chain rule, the adjoint vectors, v_n^* , satisfy the inhomogenous adjoint equation, which is an iterative equation

$$\begin{aligned} v_n^* &= (D_u f)^T(u_n, \mathcal{S}) v_{n+1}^* + \frac{1}{N} (D_u J)^T(u_n) \\ v_{N+1}^* &= 0. \end{aligned} \quad (2.9)$$

The inhomogenous adjoint equation is solved backward in time with the zero vector as the initial condition at $N + 1$. The homogeneous adjoint solution is defined by setting the source term to zero, which yields

$$q_n^* = (D_u f)^T(u_n, \mathcal{S}) q_{n+1}^*. \quad (2.10)$$

This equation, which is solved backward in time with a non-zero initial condition at N , is an iterative application of the chain rule, at a fixed \mathcal{S} , to the definition $q_n^* := (D_u(f_{N-n} \cdot q_N^*))^T(u_n)$. In other words, the homogeneous adjoint solution at time n is the sensitivity to u_n of the solution at time N projected along q_N^* , $N \geq n$. Similar to the tangent solutions, the adjoint solutions asymptotically grow exponentially, backward in time, for almost every initial condition, i.e., for large N , $q_0^*, v_0^* \sim \mathcal{O}(e^{\lambda_1 N})$. Intuitively, we can understand this growth as complementary to the growth of tan-

gent solutions. That is, since infinitesimal perturbations to the state increase in norm exponentially forward in time, we expect the solution at a given time to be more sensitive (exponentially) to its far past compared to its recent past. The Lyapunov exponents characterizing the adjoint dynamics' growth are the same as those for the tangent equations. The tangent and adjoint solutions are connected by bi-orthogonality, which means that, homogeneous adjoint solution and tangent solutions that are associated with two different Lyapunov exponents, are orthogonal to each other. In particular, the tangent solutions with negative LEs, which span E^s , must be orthogonal to adjoint solutions associated to positive or zero LEs. In other words, it can be shown that the unstable adjoint subspace, consisting of all the adjoint solutions that exponentially grow (at most at d_u different asymptotic rates) backward in time, is orthogonal to E^s and E^c . Similarly, the stable adjoint subspace is orthogonal to E^u and E^c [159].

2.2.3 Automatic differentiation to compute tangent and adjoint solutions

We consider another class of linear perturbation methods: automatic differentiation (AD). Given a program, with *output* O and *input* I , where O and I can be scalar or vector-valued, AD obtains the derivative dO/dI . In forward mode, the program is traversed sequentially, and each line is differentiated with respect to I using the derivatives of the variables computed in the previous lines. Ultimately, dO/dI is obtained exactly. In reverse mode AD, the derivative dO/dI is again obtained exactly, but by traversing the program in reverse order and using the chain rule to update the derivatives.

As noted in section 2.2.1 and section 2.2.2, each of the four linear perturbation solutions discussed (homogeneous/inhomogeneous tangent/adjoint solutions) can be written in the form of a derivative. By specifying O and I appropriately, all of them can be computed through AD. From their derivative-based definitions, the functions that specify O only require the primal solver (i.e., a time-integrator, $f(\cdot, \mathcal{S})$) and

the definition of the objective function. As a result, in order to compute linear perturbation solutions, AD does not require the user to compute the Jacobian, and eliminates the need for tangent/adjoint calculations through their respective iterative equations. For example, consider the AD solution of the inhomogeneous adjoint equation (Eq. 2.9). To compute it, we define a function that returns a value $O = (1/N) \sum_{n=0}^{N-1} J(f_n(u_0), \mathcal{S})$. The function takes the argument $I = u_0$ time-integrates the primal for N steps, evaluates J at each step for averaging at the end. Then, dO/dI computed by applying AD in reverse-mode is the solution of the inhomogeneous adjoint equation at time 0. The reader is referred to texts on AD (e.g. Ch 3 [124], Ch. 10 [122] and Ch. 15 [123] of Griewank and Walther) for the application of AD to replace tangent/adjoint solvers. In this chapter, our focus is the implementation of the shadowing algorithm [206] to compute sensitivities. We use AD to replace the tangent/adjoint solvers needed within the shadowing algorithm. The inputs and outputs to AD must be defined appropriately in the AD version of shadowing, as discussed in 2.10.

While development time is reduced by AD through the elimination of hand-differentiation, AD has compile-time and run-time overheads, which depend on the AD software used. In this chapter, we use the AD package `Zygote.jl` [145] in Julia, which uses the language’s multiple dispatch feature to compute derivatives. Several AD library options exist in languages popular in scientific computing [19]. Some combine modern language-level features (e.g. multiple dispatch or operator overloading on dynamic types in Julia) with algorithmic advances [146] to achieve time and memory efficiency [227] when compared to traditional solvers for perturbation equations, (see e.g. optimization in `PerforAD` in Python [139]).

In Figure 2-1, we plot the l^2 norms of linear perturbations computed using the four different methods discussed. The primal system that supplies f is the chaotic acoustic model (section 2.4). All of the linear perturbation methods evolve with exponentially increasing norms. As noted earlier, finite difference saturates on the order $\mathcal{O}(1/\epsilon)$ since the attractor is bounded and, therefore, so is the norm of the difference between any two solutions. The finite difference results shown as green Y’s




Figure 2-1: l^2 norms of the perturbation vectors computed through the homogeneous tangent (blue triangle), adjoint equations (orange circle), finite difference (green Y), forward-mode AD (red triangle) and reverse-mode AD (purple plus), are shown as a function of time, for the time-delayed model of section 2.4.

in Figure 2-1 are calculated with an initial perturbation with norm 10^{-4} . The norm of the finite difference increases exponentially before saturating at about 10^4 . Forward-mode AD results, which compute the tangent solutions exactly, closely approximate the latter. The tangent solutions computed using Eq. 2.4 and forward-mode AD, and the adjoint solution computed using Eq. 2.10 and reverse-mode AD, show unbounded exponential growth. The slopes of the perturbations on the logarithmic scale (≈ 0.2) indicate the largest Lyapunov exponent of the chaotic acoustic model (section 2.4). This is a manifestation of the *butterfly effect*. Next, we explain how this effect leads to the breakdown of traditional sensitivity algorithms in chaotic problems.

2.2.4 The problem with computing sensitivities of ergodic averages using conventional methods

To compute the sensitivity of a time-averaged quantity, one could potentially use the tangent, adjoint solutions, or finite difference, or AD. Consider the problem of sensitivity computation of a finite-time average $\langle J \rangle_N$. Using the tangent or adjoint

solutions solved iteratively using Eq. 2.3 and Eq. 2.9 (or computed using AD) respectively, one can obtain the derivative of $\langle J \rangle_N$ as

$$d_s \langle J \rangle_N(u_0, s) = \frac{1}{N} \sum_{n=0}^{N-1} x_n^* \cdot v_n \quad (2.11)$$

$$= \frac{1}{N} \sum_{n=0}^{N-1} x_n \cdot v_n^*, \quad (2.12)$$

where, for notational convenience, we have defined $x_n := (\partial_s f)(u_{n-1}, \mathcal{S})$, and $x_n^* := (D_u J)^T(u_n, \mathcal{S})$. Both equations, which can be derived using the chain rule, along with their AD counterparts discussed in section 2.2.3, are standard in sensitivity analysis. To reap computational benefits, Eq. 2.11, or forward-mode AD, is used when the set of observables has a larger dimension than the parameter space. The tangent solution $\{v_n\}$ corresponding to a parameter, can be used to compute the derivatives of time averages of all the observables with respect to that parameter. By contrast, when the number of parameters exceeds the number of observables, Eq. 2.12, or reverse-mode AD is the preferred approach to compute sensitivities since the same sequence $\{v_n^*\}$ computed for a given $\langle J \rangle_N$ is used to compute the gradient of $\langle J \rangle_N$ with respect to all the parameters in \mathcal{S} .

In chaotic systems, the above approach yields values of $d_s \langle J \rangle_N$ that exponentially increase with N , as $N \rightarrow \infty$. However, the quantity of interest, $d_s \langle J \rangle$ (Eq. 2.2), in which the limit $N \rightarrow \infty$ is taken before the derivative with respect to s , is bounded. Hence the derivative of the ergodic average $d_s \langle J \rangle$ is not the same as the derivative of the finite-time average $d_s \langle J \rangle_N$ in the limit $N \rightarrow \infty$. Thus, conventional methods for sensitivity computation are not applicable to the computation of derivatives of ergodic averages in chaotic systems. One early approach to circumvent this problem is the ensemble sensitivity method [167, 96] in which $d_s \langle J \rangle$ is approximated by a sample average of sensitivities computed by using Eq. 2.11 or Eq. 2.12 over a small N (comparable to one Lyapunov time). The accuracy of this method improves as N increases, provided that the number of samples increases exponentially with N . This makes the method prohibitively expensive in practice [96, 69]. In the next section,

we describe Non-Intrusive Least Squares Shadowing (NILSS) due to Ni *et al.* [206], which is a more efficient approach for computing the same quantity.

2.2.5 Non-Intrusive Least Squares Shadowing

Owing to the *shadowing lemma* (see e.g. [154]) for uniformly hyperbolic systems, it has been shown by Wang [271] that there exists a unique perturbation direction v^{sh} – the tangent shadowing perturbation – for which the tangent equation (Eq. 2.3) has a bounded solution for all time. The shadowing perturbation v^{sh} is an inhomogeneous tangent solution (i.e., v^{sh} satisfies Eq. 2.3). However, unlike the conventional tangent solution, v , it does not exhibit an unstable growth. This constraint is used by the NILSS algorithm [206] to solve for v^{sh} over a long, but finite-time, duration. In particular, NILSS [206] constructs an approximation of the shadowing perturbation by subtracting from v an unstable tangent vector at every point along a trajectory. The unstable tangent vector to be subtracted is represented in an orthonormal basis of the unstable tangent subspace (E^u). The orthonormal basis, in turn, is computed by propagating at least as many tangent vectors as the dimension of the unstable subspace, under the homogeneous tangent dynamics (Eq. 2.4), along with repeated normalization. This procedure is typically used in the computation of Lyapunov vectors [159, 113].

Let Q_n be a $d \times d_u$ matrix whose columns form an orthonormal basis of $E^u(u_n)$. In the NILSS algorithm [206], the total time duration N is divided into multiple short time segments, *checkpoints*, such that each short segment is comparable to the Lyapunov time. We shall simplify the setting by considering a time segment to be one timestep (i.e., every timestep is a checkpoint); we delay a discussion on this simplification until the end of this section. The tangent shadowing perturbation can be expressed as

$$v_n^{\text{sh}} = v_n + Q_n a_n, \quad 1 \leq n \leq N, \quad (2.13)$$

where $a_n \in \mathbb{R}^{d_u}$ is a vector of coefficients. NILSS computes the sequence of vectors $\{a_n\}$ along the trajectory by solving a minimization problem for the norms of $\{v_n^{\text{sh}}\}$. For a complete description of the NILSS algorithm, the reader is referred to [206], where the derivation of the algorithm for the time-continuous case is presented. Since perturbations along the center direction neither grow nor decay exponentially, in NILSS, the center direction is excluded from Q . Its effect on the sensitivity is added back later (see section 2 of [206]). When the map f is a time-discretized ODE, it has a center direction corresponding (but not exactly equal) to the center direction of the ODE. Thus, we also take into account, in the discrete algorithm, modifications due to this center direction, when f models ODEs, as we shall see in section 2.3.4.

The NILSS problem minimizes the norms of the shadowing perturbation sequence $\{v_n^{\text{sh}}\}$. The Lagrangian of this optimization problem can be written as

$$\mathcal{L}^{\text{sh}}(\{a_n\}, \{\beta_n\}) := \sum_{n=1}^N \|v_n^{\text{sh}}\|^2 + \sum_{n=1}^{N-1} \beta_n (a_{n+1} - R_{n+1}a_n - \pi_n) \quad (2.14)$$

$$= \sum_{n=1}^N \left(\|v_n\|^2 + \|a_n\|^2 + 2v_n \cdot Q_n a_n \right) + \sum_{n=1}^{N-1} \beta_n (a_{n+1} - R_{n+1}a_n - \pi_n), \quad (2.15)$$

where Eq. 2.15 uses Eq. 2.13, and the fact that $Q_n^T Q_n$ is the $d_u \times d_u$ identity matrix. Here, $\{\beta_n\}$ is a sequence of Lagrange multipliers that imposes a sequence of equality constraints at every timestep to ensure the continuity of the shadowing perturbation (section 2.3.3). We solve the above problem to obtain a sequence $\{a_n\}$, and then, to obtain the shadowing perturbation through Eq. 2.13. Subsequently, we compute the required sensitivity through Eq. 2.11, with the (exponentially growing) tangent solution, v_n , replaced with the shadowing tangent solution v_n^{sh} . This yields

$$d_s \langle J \rangle \approx \frac{1}{N} \sum_{n=0}^{N-1} x_n^* \cdot v_n^{\text{sh}}, \quad (2.16)$$

as shown in Appendix C of [206], or in Theorem LSS of [271]. The same shadowing perturbation v_n^{sh} is used to compute the sensitivities with respect to all $J \in \mathcal{J}$. On

the other hand, the tangent NILSS algorithm has to be repeated for every parameter $s \in \mathcal{S}$ in order to compute the corresponding shadowing perturbations. When the parameter space is higher-dimensional when compared to the observable space, the adjoint version of NILSS is preferred.

Adjoint non-intrusive least squares shadowing

While the adjoint algorithm can be obtained via reverse-mode automatic differentiation of tangent NILSS, the theoretical basis for the existence of an adjoint shadowing perturbation is developed in [207]. The adjoint algorithm, known as Non-Intrusive Least Squares Adjoint Shadowing (NILSAS), is presented in [211] with an application to a fluid flow problem in [208]. Here, we focus on the discrete time case, and for simplicity, each time segment corresponds to one iteration of the map f . Analogous to tangent NILSS, in the adjoint version, an adjoint shadowing perturbation $v_n^{\text{sh}*}$ is computed, which solves the inhomogeneous adjoint equation (Eq. 2.9). Extending the analogy further, an unstable adjoint vector is subtracted from the conventional adjoint solution to obtain $v_n^{\text{sh}*}$ as

$$v_n^{\text{sh}*} = v_n^* + Q_n^* a_n^*, \quad (2.17)$$

where Q_n^* is an orthonormal basis for the unstable adjoint subspace, $(E^s \oplus E^c)^\perp$, and $a_n^* \in \mathbb{R}^{d_u}$ is a set of coefficients. The orthonormal basis $Q_n^* \in \mathbb{R}^{d \times d_u}$ is achieved by iterating at least d_u adjoint vectors backward in time, using Eq. 2.10, repeatedly normalizing with QR factorization. That is, the orthonormalization procedure is identical to that for $\{Q_n\}$, but using the sequence of Jacobian transposes, instead of the Jacobians, and with time-reversal. The particular set of coefficients a_n^* needed to find the adjoint shadowing sequence $v_n^{\text{sh}*}$, which is a bounded solution of the inhomogeneous adjoint equation, is found as a solution of a least squares problem, which is also analogous to Eq. 2.15. In order to compute the quantity of interest, we replace the conventional (exponentially growing) inhomogeneous adjoint solution in Eq. 2.12

with the adjoint shadowing solution

$$d_s \langle J \rangle \approx \frac{1}{N} \sum_{n=0}^{N-1} v_n^{\text{sh}*} \cdot x_n. \quad (2.18)$$

We now comment on the error-vs-cost trade-off of the tangent and adjoint shadowing algorithms, noting that a careful analysis of this trade-off is problem-specific and beyond the scope of this chapter. In both algorithms, the most expensive computation in the shadowing algorithm is the solution of the minimization problem. The size of the minimization problem is directly proportional to d_u and N . For the same overall time duration N , choosing a larger time segment between checkpoints reduces the size of the minimization problem, since the QR factorization and the equality constraints that enter the problem (Eq. 2.15) are only computed at the boundaries of the time segments.

For example, in the time delay acoustic problem of this chapter, we could theoretically choose a segment size over which the linear perturbations increase in norm by, say, a factor of 2. From Figure 2-1, such a segment size is about 1 time unit (100 timesteps). Then, we need to perform QR factorization only every 100 timesteps. While the QR factorization in itself is not the computational bottleneck, the size of the least squares problem shrinks by a factor of 100, as compared to checkpointing every timestep. In spite of the additional cost, we choose to checkpoint every timestep in the considered problem. However, we choose a size N such that the $\mathcal{O}((Nd_u)^3)$ calculation of the least squares solution is neither memory-constrained nor is a prohibitive computational expense. Then, we repeat the shadowing algorithm M times and sample-average the shadowing sensitivities. This procedure is effectively the same as computing the sensitivity of an MN -time average by executing the shadowing algorithm once, provided that N is large enough for the convergence of the ergodic averages. The reason for sample-averaging shadowing sensitivities as opposed to segmenting a long shadowing algorithm is that we observe, for the time-delayed acoustic model, higher condition numbers of the least squares problem. Thus, we take the approach of checkpointing every timestep, and solving smaller least squares problems.

The size of each problem is chosen large enough for ergodic averages to converge while also curtailing the computed shadowing perturbations (from the solution of the least squares problem) to an $\mathcal{O}(1)$ norm, at every timestep. The advantage of this approach is the simpler program whose computational cost is nearly the same as the MN -sized checkpointed NILSS, but produces better-behaved shadowing perturbations.

2.2.6 Error in shadowing and alternatives

The NILSS algorithm is not guaranteed to converge to the true value of the sensitivity, as $N \rightarrow \infty$. To see why, we first note that the shadowing sensitivities computed by NILSS, which are the right hand sides of Eq. 2.16 and Eq. 2.18, are ergodic averages along a true orbit of a system at an infinitesimally perturbed s . For a mathematically rigorous explanation, the reader is referred to [271]; a qualitative explanation is also included in section 2.6. Now, due to a perturbation in s , the stationary probability distribution on the attractor is perturbed as well, but this perturbation is excluded by NILSS. In general, the ergodic averages along a shadowing orbit (a true orbit) of a perturbed system do not converge to the expectation with respect to the stationary distribution of the unperturbed system. Hence, there is a systematic error in NILSS, which has been recently studied in [209], along with the corrections that can be made to reduce the error. In view of this shadowing error, we must mention that a few alternatives have recently appeared. In particular, the space-split sensitivity method (Chapter 7) is an ergodic-averaging method to compute Ruelle’s linear response formula [233], which specifies the required sensitivity exactly. However, based on the formulation in this thesis [62, 247], it is more complex to implement than the shadowing method. Thus, in applications where a systematic error in the computed sensitivities is not a serious impairment, such as in the parameter estimation and data assimilation problems considered in sections 2.5 and 2.6, respectively, shadowing methods may be preferred. Another sensitivity computation method, also based on Ruelle’s linear response formula [233], is known as blended response [1], in which short-term and long-term responses are computed using different methods. For the long-term response to unstable perturbations, Ruelle’s formula, which is exact, is

approximated using a Fluctuation-Dissipation theorem for non-equilibrium settings [185]. However, this approach is computationally expensive and also inexact. Thus, we focus on the shadowing-based methods for computing sensitivities in this chapter. Moreover, shadowing methods have successfully been applied to dissipative models in fluid mechanics [208, 141], as relevant to this study.

2.3 Tangent/adjoint shadowing algorithm

We provide a step-by-step description of the tangent and adjoint NILSS algorithms. The reader is referred to [206] and [211] for the original descriptions of tangent and adjoint NILSS, respectively. We consider the discrete algorithm without the checkpointing scheme. We also adopt a simplified presentation for which the same program can be used for implementing both tangent and adjoint NILSS, with a minimal modification. Further, here, we introduce AD to compute the needed tangent and adjoint solutions. This automatic-differentiated unified program for the tangent/adjoint NILSS, shall be referred to as the AD shadowing algorithm.

2.3.1 Shadowing algorithm: inputs and outputs

We present a shadowing algorithm that constructs a sequence of shadowing perturbations $v_n^{\text{sh}}, 1 \leq n \leq N$. The algorithm takes as inputs, a sequence of $d \times d$ matrices A_n , and a sequence of d -length vectors, b_n , to return tangent or adjoint shadowing sensitivities (defined in Eq. 2.16 and Eq. 2.18, respectively) in the following two scenarios.

- **Case 1 (tangent):** The sequence $\{A_n\}$ is set to the Jacobian matrix sequence $\{D_u f(u_n, \mathcal{S})\}$ along a reference trajectory $\{u_n\}$. The sequence b_n is the parameter perturbation at n , i.e., $b_n = x_n$. Then, the algorithm returns the sequence of tangent shadowing perturbation vectors at u_1, u_2, \dots, u_N , namely, $v_1^{\text{sh}}, v_2^{\text{sh}}, \dots, v_N^{\text{sh}}$. These shadowing perturbations can be used to approximately compute the l sensitivities, $d\langle \mathcal{J} \rangle / ds$.

- **Case 2 (adjoint):** Now, on the other hand, defining $n' := N + 1 - n$ suppose A_n is the transpose of the Jacobian matrix at n' , that is, $A_n = (D_u f)^T(u_{n'})$. Set $b_n := x_{n'+1}^*$. In this case, the algorithm returns the sequence of adjoint shadowing perturbation vectors at $u_{N+1}, u_N, u_{N-1}, \dots, u_2$, namely, $v_1^{\text{sh}}, v_2^{\text{sh}}, \dots, v_N^{\text{sh}}$ (in the unified presentation, we drop the superscript “*” used for adjoint solutions). The adjoint shadowing perturbation sequence is used to compute the p sensitivities, $D_S \langle J \rangle$.

A time reversal is accomplished in adjoint shadowing simply by reversing the indexing of the input sequences; the output sequence of adjoint shadowing perturbations is obtained in time-reversed order. We use the term *shadowing perturbation* to refer to both tangent and adjoint shadowing perturbations. Note that u_0 must be a point on the attractor sampled according to the stationary probability distribution on the attractor. That is, the primal system must be simulated for a run-up time long enough for time-averages to converge. A solution state obtained after such a run-up time is chosen as u_0 .

2.3.2 Evolution of homogeneous and inhomogeneous perturbations with repeated normalization

Our goal is to compute the tangent or adjoint shadowing perturbation using Eqs. 2.13 and 2.17, respectively. Toward this goal, we solve i) at least d_u homogeneous tangent equations, or d_u homogeneous adjoint equations, and ii) p inhomogeneous tangent equations or l inhomogeneous adjoint equations. The common form of the homogeneous equation, which amounts to solving the tangent equation forward (in case 1) or the adjoint equation backward in time (in case 2), is given by

$$q_n^i = A_{n-1} q_{n-1}^i, \quad n = 1, \dots, N, \quad 1 \leq i \leq d_u. \quad (2.19)$$

We define Q_n to be an $n \times d_u$ matrix with columns q_n^i . The following equation gives the evolution of the inhomogeneous tangent solution forward in time in case 1, and

the inhomogeneous adjoint solution backward in time in case 2,

$$v_n = A_{n-1}v_{n-1} + b_n, \quad n = 1, \dots, N. \quad (2.20)$$

At each n , we normalize both the homogeneous and inhomogeneous perturbations by QR factorization. We choose d_u pseudo-random vectors in \mathbb{R}^d as initial conditions q_0^i , $1 \leq i \leq d_u$. The initial condition for Eq. 2.20, v_0 , is set to $0 \in \mathbb{R}^d$. Beginning with $n = 1$, we perform the following loop until $n = N$.

1. Advance Eq. 2.19 by one timestep for each $1 \leq i \leq d_u$. This can be written as $Q_n \longleftarrow A_{n-1}Q_{n-1}$.
2. QR-factorize the matrix Q_n and set Q_n to the obtained “ Q ”. Let the “ R ” from QR factorization be stored as R_n . Thus, each q_n^i , $1 \leq i \leq d_u$ is now a unit vector.
3. Obtain v_n from v_{n-1} by advancing Eq. 2.20 by one timestep.
4. Set $\pi_n := Q_n^T v_n$, which is a d_u -length vector of orthogonal projections of v_n along q_n^i .
5. Project out the unstable components of v_n . That is, set $v_n \longleftarrow v_n - \pi_n Q_n$.
6. Go to step 1 with $n \longleftarrow n + 1$, or stop if $n = N$.

First, using the above orthonormalization procedure, Q_n converges to an orthonormal basis for the unstable tangent (adjoint) subspace in case 1 (case 2). Secondly, in case 1, we note that the above n -loop must be executed only once for the sensitivity with respect to s of all $J \in \mathcal{J}$. Similarly, in case 2, the n -loop must be called just once if we wish to compute the sensitivity of $\langle J \rangle$ with respect to all the parameters \mathcal{S} . In other words, to obtain $D_{\mathcal{S}}\mathcal{J}$, in case 1 (tangent shadowing), the n -loop must be run as many times as the number of parameters ($= p$), and in case 2 (adjoint shadowing), as many times as the number of objective functions ($= l$). Thirdly, the sequence of matrices R_n can be used to obtain the Lyapunov exponents. In particular, if the k th

diagonal element of the matrix R_n is written as R_n^k , then, the k th Lyapunov exponent $\lambda_k \approx (1/N) \sum_{n=0}^{N-1} \log |R_n^k|$. This can be easily seen by recasting the definition of Lyapunov exponents ([15]) and as an ergodic average.

2.3.3 Minimizing the growth of the shadowing perturbation sequence

At the end of the n -loop described in section 2.3.2, we have at our disposal the following sequences of vectors or matrices, where at each $n \leq N + 1$,

- v_n is the inhomogeneous perturbation orthogonalized with respect to the unstable tangent (adjoint) subspace in case 1 (case 2).
- π_n consists of the orthogonal projections (before the orthogonalization) of v_n on the unstable tangent (adjoint) subspace in case 1 (case 2).
- Q_n is a $d \times d_u$ matrix that forms an orthonormal basis for the unstable tangent (adjoint) subspace at each n (n') in case 1 (case 2), and,
- R_n is a $d_u \times d_u$ matrix that contains the one-step growth factors of Q_n under the tangent (adjoint) dynamics in case 1 (case 2).

In practice, a finite spin-up time, typically on the order of Lyapunov time, is needed for the convergence of Q_n to an orthonormal basis for the true unstable (tangent/adjoint) subspace. We can write the ansatz for the shadowing perturbation sequence (Eq. 2.13 and Eq. 2.17) in a form that is applicable to both tangent and adjoint shadowing sequences, denoted here as v^{sh} ,

$$v_n^{\text{sh}} = v_n + Q_n a_n. \tag{2.21}$$

Here the sequence a_n is the unknown d_u -length vector, which we shall solve for. In case 2, the sequence a_n , and subsequently v_n^{sh} , are obtained in time-reversed order by virtue of time-reversing the inputs A_n, b_n to the n -loop. In particular, v_n^{sh} is the

adjoint shadowing perturbation at time $N + 2 - n$. In order to solve for a_n , we start by multiplying Eq. 2.21 by A_n and adding b_{n+1} to both sides of the equation,

$$A_n v_n^{\text{sh}} + b_{n+1} = A_n v_n + b_{n+1} + A_n Q_n a_n. \quad (2.22)$$

Since the shadowing perturbation solves Eq. 2.20, the left hand side is v_{n+1}^{sh} . Using steps 3 to 5 of the n -loop in section 2.3.2, the first two terms on the right hand side of Eq. 2.22 become $v_{n+1} + Q_{n+1}\pi_{n+1}$. From step 2 of the n -loop, $A_n Q_n = Q_{n+1}R_{n+1}$. Thus,

$$v_{n+1}^{\text{sh}} = v_{n+1} + Q_{n+1}\pi_{n+1} + Q_{n+1}R_{n+1}a_n. \quad (2.23)$$

From Eq. 2.21, the left hand side of the above equation is also equal to $v_{n+1} + Q_{n+1}a_{n+1}$. We obtain the following iterative relationship for a_n , after multiplying both sides by Q_{n+1}^T

$$a_{n+1} = \pi_{n+1} + R_{n+1}a_n. \quad (2.24)$$

This is the equality constraint that must be added to the NILSS problem, whose Lagrangian is in Eq. 2.15. Hence, Eq. 2.24 is also one of the KKT conditions ($D_{\beta_n} \mathcal{L}^{\text{sh}} = 0$) of the NILSS optimization problem. Although one can theoretically solve Eq. 2.24 starting from a random guess for a_1 and iterating, this does not provide accurate results in practice; Eq. 2.24 is not a well-conditioned problem for $\{a_n\}$ due to the exponential growth of the round-off errors in $\{R_n\}$, which tend to accumulate upon iteration. Thus, following [206, 211], we resort to the direct method of solving for the entire sequence $\{a_n\}$ at once (Appendix A of [211] and [206]). The direct method is to solve the following system of linear equations for $\{a_n\}$

$$GX = H, \quad (2.25)$$

where

- G is an $Nd_u \times (N + 1)d_u$ block matrix with $d_u \times d_u$ blocks given by

$$G := \begin{bmatrix} -R_1 & I & 0 & \cdots & 0 & 0 \\ 0 & -R_2 & I & \cdots & 0 & 0 \\ 0 & \cdots & \cdots & \cdots & 0 & 0 \\ \cdots & \cdots & \cdots & \cdots & \cdots & \cdots \\ 0 & \cdots & \cdots & -R_{N-1} & I & 0 \\ 0 & \cdots & \cdots & \cdots & -R_N & I \end{bmatrix}, \quad (2.26)$$

where I is the $d_u \times d_u$ identity matrix,

- X is an $N \times d_u$ vector consisting of $[a_0, \cdots, a_N]$ and,
- H is an Nd_u -length vector containing the sequence $[\pi_1, \cdots, \pi_N]$.

The solution of the underdetermined system that minimizes the norm of X is given by

$$X = G^T(GG^T)^{-1}H.$$

2.3.4 Modifications due to the center direction

Whether in tangent or adjoint shadowing, a better accuracy is obtained if the center direction, which is approximately the right hand of an ODE when f is a time-discretization of the ODE, is given a special treatment. In tangent shadowing, d_u can be set to the number of positive LEs plus 1, so that the center direction is treated as an unstable direction. However, in some problems, this may increase the condition number of the least squares problem for X . This leads to a poorer minimization of X , which in turn increases the norm of the shadowing perturbation, when compared to the following alternative. As suggested in [206], we project out the center components of both the homogeneous and inhomogeneous tangents and add the contribution to the sensitivity due to the center perturbation, in the final step. We discuss the modification to which this leads in the n -loop (section 2.3.2). Then, we discuss the modification in the calculation of the sensitivity in the next subsection. In

the n -loop, in addition to step 2, we must also subtract from Q_n its projection along \mathcal{F} , which is approximately the center direction, as: $Q_n \leftarrow Q_n - \mathcal{F}_n \mathcal{F}_n^T Q_n / \|\mathcal{F}_n\|^2$, where $\mathcal{F}_n := \mathcal{F}(u_n, \mathcal{S})$. Similarly, for v_n , after step 3, $v_n \leftarrow v_n - \mathcal{F}_n \mathcal{F}_n^T v_n / \|\mathcal{F}_n\|^2$. Next, we discuss modifications to adjoint shadowing due to the center perturbation. As we noted earlier, the tangent and adjoint subspaces corresponding to two different Lyapunov exponents are perpendicular to each other (see e.g. Appendix B in [207] for a proof; we also show numerical results verifying this fact for the time-delayed system considered, in Figure 2-6). This orthogonality gives rise to the constraint (derived in section 5.4 of [207]):

$$(1/N) \sum_{n=1}^N (v_n^{\text{sh}})^T \mathcal{F}_n = 0.$$

By definition, since $v_n^{\text{sh}} = v_n + Q_n a_n$, this leads to,

$$\sum_{n=1}^N \left(v_n^T \mathcal{F}_n + a_n^T Q_n^T \mathcal{F}_n \right) = 0. \quad (2.27)$$

This condition leads to one more equation (adding one more row to G) while solving the NILSS problem (section 2.3.3). The n -loop in adjoint shadowing need not be modified.

2.3.5 Computation of the sensitivities

Having obtained the sequences $\{a_n\}$, $\{v_n\}$ and $\{Q_n\}$, the shadowing perturbation is determined, for $1 \leq n \leq N$, as

$$v_n^{\text{sh}} = v_n + Q_n a_n. \quad (2.28)$$

With the shadowing perturbation, the sensitivities can be computed as though the system were not chaotic (Eq. 2.11 and Eq. 2.12). That is, in adjoint shadowing,

$$d_s \langle J \rangle = \frac{1}{N} \sum_{n=1}^N v_n^{\text{sh}} \cdot x_{n'+1}. \quad (2.29)$$

In the case of tangent shadowing, we add to the sensitivity in Eq. 2.11, the contribution from the center direction, if treating the center direction separately as described in section 2.3.4,

$$d_s \langle J \rangle = \frac{1}{N} \left(\sum_{n=1}^N x_n^* \cdot v_n^{\text{sh}} + \frac{v_n^T \mathcal{F}_n}{\|\mathcal{F}_n\|^2} (J_n - \langle J \rangle_N) \right), \quad (2.30)$$

where J_n in the above equation is the objective function at u_n . In both tangent and adjoint shadowing, the projections onto the center direction, $v_n^T \mathcal{F}_n$ and $Q_n^T \mathcal{F}_n$, are stored during the n -loop. However, the n -loop needs to be modified to account for the center direction only while performing tangent shadowing.

2.3.6 Automatic differentiation

As discussed in section 2.2.3, we can replace tangent/adjoint solvers with forward/reverse-mode automatic differentiation, respectively. In the n -loop (section 2.3.2), we can introduce AD to advance Q_n, v_n . Hence, AD-shadowing only requires the primal solver to be supplied by the user, as opposed to primal, tangent and adjoint solvers. The AD-version of the n -loop is shown in 2.10.

We remark that for the AD version of the shadowing algorithm, an exploration of various techniques for memory and time-efficiency of AD [35, 144], such as combining primal-tangent/adjoint solver, is needed. These approaches may lead to taking longer timesteps without compromising on accuracy by utilizing the fact that AD is an exact method, which does not increase the numerical error in the perturbations. This more involved approach to AD shadowing, must be numerically investigated for a given problem to determine whether (or not) it leads to a realizable computational advantage (due to AD overheads) in practice.

2.4 The time delayed model for thermoacoustics

Chaotic thermoacoustic oscillations originate from two main physical nonlinearities, which are deterministic. First, the heat released by the flame is a nonlinear function of

the acoustic perturbations at the flame's base, i.e. the flame saturates nonlinearly [87, 88]. Both experimental investigations [116, 149, 117, 150] and numerical studies [274, 153, 214] showed that the nonlinear flame saturation may cause a periodic acoustic oscillation to become chaotic, by either period doubling, or Ruelle-Takens-Newhouse, or intermittency scenarios [205, 204], which are common in fluid dynamic systems [90, 200, 91]. The numerical studies of [274, 153, 214] showed that the nonlinear flame saturation may generate chaotic acoustic oscillations even in laminar flame models, where the turbulent hydrodynamics is not modelled. We introduce a nonlinear time-delayed model of chaotic thermoacoustic instabilities. We demonstrate that shadowing obtains useful sensitivities of this model, in the chaotic regime. We begin by describing the flame duct model of combustion in a horizontal Rijke tube [253], open to the atmosphere on both ends. The inviscid momentum and energy equations are linearized about the mean flow to yield,

$$\frac{\partial u}{\partial t} + \frac{\partial p}{\partial x} = 0 \quad (2.31)$$

$$\frac{\partial p}{\partial t} + \frac{\partial u}{\partial x} + \zeta p - \dot{q} \delta(x - x_f) = 0, \quad (2.32)$$

where, $u(x, t)$ and $p(x, t)$ are the acoustic velocity and pressure at the one-dimensional spatial location x and at time t . The pointwise heat-release source is $\dot{q} \delta(x - x_f)$, where $\delta(x - x_f)$ is the Dirac delta centered at x_f . The constant ζ is a parameter that models acoustic damping [163, 162]. We consider a Galerkin modal decomposition in a Fourier basis [253] of the acoustic velocity and pressure fields, which transforms Eqs. 2.31-2.32 into a set of time-delayed coupled oscillators

$$\begin{aligned} \frac{d\eta_j}{dt} - j \pi \theta_j &= 0, \\ \frac{d\theta_j}{dt} + j \pi \eta_j + \zeta_j \theta_j + 2\beta \dot{q}(u_f(t - \tau)) \sin(j\pi x_f) &= 0, \end{aligned} \quad (2.33)$$

where $j = 1, \dots, d_g$, indicates the index of the Galerkin modes, with d_g being their total number,

$$u_f(t) = \sum_{k=1}^{N_g} \eta_k(t) \cos(k\pi x_f), \text{ and} \quad (2.34)$$

$$\zeta_j = c_1 j^2 + c_2 j^{1/2}. \quad (2.35)$$

In the above system of equations, η_j indicate the velocity modes and θ_j indicate the pressure modes. Modal damping is represented by ζ_j and \dot{q} is the rate of heat release at the flame location x_f . The function $\dot{q}(u) = \sqrt{|1.0 + u|} - 1$ is a modified King's law [253], which is non-differentiable at $u = -1$. Along trajectories of this system, we may encounter states corresponding to $u_f = -1$, where the Jacobian does not exist. We need to avoid this non-differentiability in order to compute the tangent/adjoint/AD solutions needed for the shadowing algorithms. Thus, we follow the approach taken in [141], wherein the function \dot{q} is approximated by a polynomial around the non-differential point. In particular, when $-1.01 \leq u \leq -0.99$, we take

$$\dot{q}(u) = -1 + 1750 (1 + u)^2 - 7.5 \times 10^6 (1 + u)^4,$$

where the coefficients have been estimated by regression. The flame velocity u_f affects the pressure field modes θ_j after a time delay given by a constant, τ . This models the fact that the disturbances in the flame velocity at the flame base require a finite time to traverse the flame and cause a perturbation in the heat released [172, 188].

2.4.1 Replacing the time delay with an advection equation

When solving the above system numerically, if the time delay is modelled by converting the delay parameter τ into an integer number of timesteps – that is, τ is converted into a discrete parameter from a continuous one – the state cannot be differentiated with respect to τ . We resolve this problem in order to ensure that AD/tangent/adjoint solvers are applicable, by augmenting the primal system with an

auxiliary linear advection model [141]

$$\tau \frac{\partial v}{\partial t} + 2 \frac{\partial v}{\partial y} = 0, \quad -1 \leq y \leq 1 \quad (2.36)$$

$$v(y = -1, t) = u_f(t). \quad (2.37)$$

The exact solution of the above advection equation at the right boundary is $v(y = 1, t) = u_f(t - \tau)$. The advection solution, $v(y = 1, t)$ can be used in place of the heat release model, which in turn influences the pressure modes as per Eq. 2.33. Thus, we mathematically make the overall primal system memory-less, by including the advection subsystem in Eq. 2.36 in the primal system (Eq. 2.33). The above discretized equation retains the chain rule dependence on τ and, hence, can be differentiated with respect to τ through AD. We use a Chebyshev spectral collocation method ([261] Ch. 6) to solve the advection equation. The additional cost per timestep incurred due to adding this advection subsystem (i.e., adding Eq. 2.36 to the primal set of ODEs in Eq. 2.33), depends on the spatial scale (in y) of the numerical discretization of the advection system, i.e., the number of Chebyshev collocation points. We seek to minimize the number of Chebyshev points in order keep the overall dimension of the system as small as possible. The timestep size can also be commensurately increased, ensuring the CFL condition, on decreasing the spatial resolution, which is cost-effective for computing long-time averages. With these considerations, we choose $d_c = 10$ Chebyshev points in the interval $-1 \leq y \leq 1$ and the timestep of the primal system is chosen to be $\tau/(2N_c)$. Choosing $d_g = 10$, the primal system of dimension $d = 2d_g + d_c = 30$ is time-evolved by integrating Eq. 2.33 and Eq. 2.36 using the Tsitouras Runge-Kutta time-integrator (Tsit5()) offered by the Julia package OrdinaryDiffEq [226, 225]. We remark that this auxiliary equation approach can be used in any general system with a constant time-delay in order to maintain its differentiability through AD with respect to the delay parameter.

2.4.2 Types of solutions over a range of the heat-release parameter

We fix the damping coefficients at $c_1 = 0.06$, $c_2 = 0.01$, the delay parameter at $\tau = 0.2$ [190], and numerically solve the primal system. In this section, we study the effect of the heat-release parameter β on the *type* of primal solution observed. Here, “type” refers to the three different possibilities for the asymptotic behavior of a nonlinear dissipative dynamical system, apart from convergence to a fixed point: convergence to a (a) limit cycle (periodic behavior), (b) quasiperiodic attractor and (c) chaotic attractor. These different regimes are all observed upon varying the parameter β from about 2 to 9; smaller values of $\beta \lesssim 0.8$ lead to a fixed point solution. We show these regimes as a function of β in the bifurcation diagram of Figure 2-2.

Ergodic average of acoustic energy

The bifurcation diagram in Figure 2-2 shows the acoustic energy upon time-averaging over a long time window against the heat release parameter β . This quantity is commonly used as an objective function for optimization problems in thermoacoustics, and shall therefore be used to demonstrate the discrete shadowing algorithm. We denote the acoustic energy by J_{ac} , and its ergodic/ensemble average by $\langle J_{ac} \rangle$, the latter quantity being computed numerically by time-averaging over a *long* trajectory. The instantaneous acoustic energy is the sum of the acoustic kinetic and potential energies, i.e., it is the Hamiltonian (constant of motion) of the natural acoustic system. Using Parseval’s theorem, the acoustic energy is related to the Galerkin modes. That is, the acoustic energy is defined as

$$J_{ac} = \frac{1}{2}(p^2 + u^2) = \frac{1}{4} \sum_{j=1}^{N_g} (\eta_j^2 + \theta_j^2). \quad (2.38)$$

The length of the averaging window is chosen to be the time taken for the standard error in the empirical mean values to be within 1% of the mean, when computed in the chaotic regime. This time, the time-average over which approximates the infinite-

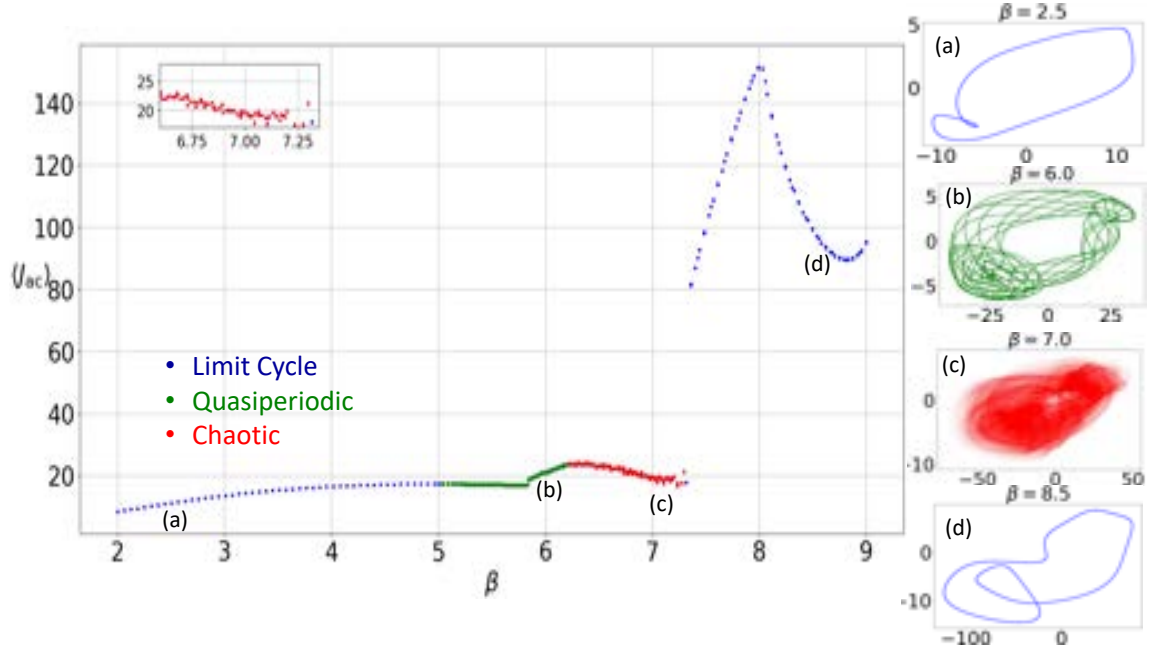


Figure 2-2: Bifurcation diagram. The values of $\langle J_{ac} \rangle$, the time-averaged acoustic energy, are color-coded according to the type of solution: periodic (blue), quasiperiodic (green) and chaotic (red). The attractors represented on the u_f - \dot{q} plane, at four different β values – $\beta = 2.5, 6.0, 7.0, 8.5$ (from top to bottom) – are shown to the right. The attractors shown on the u_f - \dot{q} plane, are also color-coded according to their type. On the top-left inset figure is shown a zoomed-in plot of $\langle J_{ac} \rangle$ -vs.- β at $\beta = 7.0$.

time/ensemble average of a given function, was determined to be about 200 time units (or 20000 timesteps, with a fixed timestep size of 0.01).

Another common objective function is the Rayleigh index, whose long-term behavior is also the subject of the sensitivity studies in this chapter. The Rayleigh index is defined as [193, 141]

$$J_{\text{ray}} := p(x_f, t) \dot{q}(t) = \frac{1}{2} \sum_{j=1}^{N_g} \zeta_j \theta_j^2. \quad (2.39)$$

The physical significance of both these objective functions is discussed later in section 2.4.3.

Limit cycles and quasiperiodicity

From Figure 9, it can be seen that $\langle J_{ac} \rangle$ increases continuously over a range of β values from 2 to about 5. In this range, the primal solution is a limit cycle, which has the effect that the values of $\langle J_{ac} \rangle$ (shown as blue dots) appear to be perfectly observed, without any noise. With $d_g = 10$ Galerkin modes and $d_c = 10$ Chebyshev points, the periodic attractor lives in a 30-dimensional space. For visualization, we show 2D phase diagrams on the $u_f-\dot{q}$ plane, for the different solution regimes. The limit cycle phase diagram is shown in blue in the top-right of the bifurcation diagram.

When $\beta \gtrsim 5$, the limit cycle transitions into quasiperiodic oscillations, which are aperiodic but appear to be *almost* periodic. For example, an iterative process of rotation on a complex unit circle (or more generally, on the surface of a d -dimensional torus) by a constant rational angle is periodic, while a rotation by a constant irrational angle, is quasiperiodic. Mathematically, a quasiperiodic solution is distinguished from a periodic solution by the number of zero Lyapunov exponents: quasiperiodic solutions have more than one while periodic solutions have exactly one zero Lyapunov exponent. We compute the Lyapunov exponents numerically (using a standard algorithm as explained in section 2.3.2) in order to classify the different types of solutions [141].

The quasiperiodic phase diagram, and the $\langle J_{ac} \rangle$ values in this regime are color-coded green in Figure 2-2. Quasiperiodicity occurs in the transition from periodic behavior to chaotic behavior. Both the periodic and quasiperiodic case are nonlinearly stable, i.e., the nonzero Lyapunov exponents are negative. This means that an applied (infinitesimal) perturbation does not grow exponentially (it may have subexponential growth) in either case.

The chaotic regime

When $6.4 \lesssim \beta \lesssim 7.3$, the solutions exhibit at least one positive Lyapunov exponent: this is the chaotic regime. In the bifurcation diagram (Figure 2-2), this regime is shown in red. The phase portrait on the $u_f-\dot{q}$ plane shows, as expected, a fractal




Figure 2-3: The first 20 LEs at $\beta = 7.0$ and $\tau = 0.2$. The QR factorization is performed every timestep, ie, at a segment length of 0.01. The total integration time considered is 200 time units. Inset: the first 3 LEs at $\beta = 7.0$ and $\tau = 0.2$. The values obtained are $\lambda_1 \approx 0.19$, $\lambda_2 \approx 0.05$, and $\lambda_3 \approx -0.07$.

attractor. The values of $\langle J_{ac} \rangle$ also appear to be erratic, revealing the presence of statistical noise due to a finite time-averaging window. In Figure 2-3, we show the first 20 Lyapunov exponents at $\beta = 7$. The value of the first exponent is about 0.2. As mentioned in section 2.2.1, the time-derivative of the ODE is a center perturbation with a zero Lyapunov exponent. The second Lyapunov exponent is about 0.05 corresponding to this center direction (it would converge to zero as the averaging time approaches infinity).

2.4.3 Acoustic energy and Rayleigh criterion as objective functions

We analyze the chaotic thermoacoustic oscillation of the primal system by studying the sensitivities of the long-time averages $\langle J_{ac} \rangle$ and $\langle J_{ray} \rangle$. Before we compute the sensitivities, we motivate our particular choice of objective functions, among many available candidates for norms [75, 109], semi-norms [187, 42], and physical measurements in this multi-physical system. For thermoacoustic systems with negligible mean flow, which cannot advect flow inhomogeneities like entropy spots, the acoustic energy and Rayleigh criterion are two suitable quantities of interest [193].

Since J_{ac} is (half) the Euclidean semi-norm of the thermoacoustic system, we are interested in calculating the sensitivity of its time average, $\langle J_{ac} \rangle$, in the interests of reducing the amplitude of chaotic oscillations. The Rayleigh index can be derived by (i) multiplying the acoustic momentum equation (2.31) by u ; (ii) multiplying the acoustic energy equation (2.32) by p ; (iii) adding them up; and (iv) integrating in the space domain. This procedure yields an equation for the evolution of the acoustic energy

$$\frac{dJ_{ac}}{dt} = - \int_0^1 \zeta p^2 dx + p_f \dot{q}, \quad (2.40)$$

where $p_f(t) := p(x_f, t)$ is the pressure at the heat source. Defining the Rayleigh index as

$$J_{ray} := p_f \dot{q}(t), \quad (2.41)$$

upon numerical discretization, we obtain

$$J_{ray} = -\dot{q}(v(y=1, t)) \sum_{j=1}^{N_g} \theta_j(t) \sin(j\pi x_f). \quad (2.42)$$

The Rayleigh index is an important cost functional that determines the stability of acoustic oscillations fed by a heat source. Physically, Eq. 2.40 states that the acoustic energy grows in time when the pressure at the heat source is sufficiently in phase with the heat release rate to exceed damping mechanisms. The acoustic energy grows up to nonlinear saturation, after which the self-sustained acoustic oscillation persists. This mechanism is commonly studied through the Rayleigh criterion [178] for the production of thermoacoustic oscillations. In chaotic oscillations, we are interested in calculating the sensitivity of the time-averaged Rayleigh index, $\langle J_{ray} \rangle$. Applying the

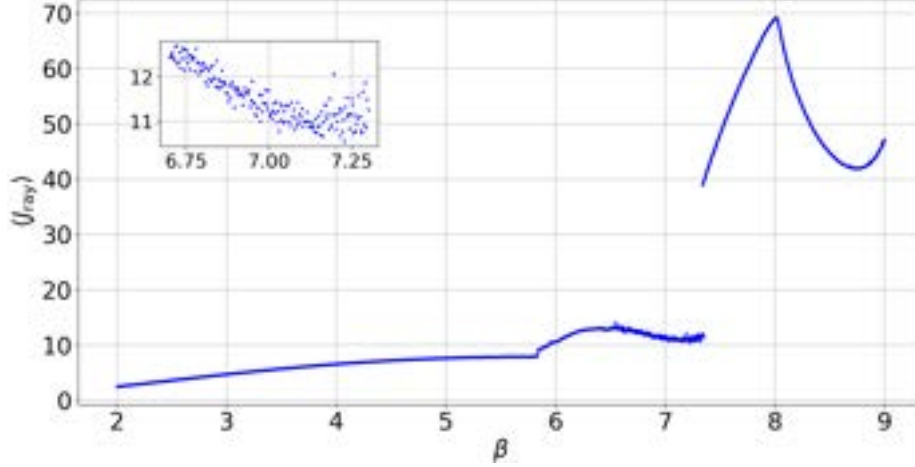


Figure 2-4: Ergodic average of the Rayleigh index, $\langle J \rangle_{\text{ray}}$. Inset: zoomed-in plot of $\langle J \rangle_{\text{ray}}$ -vs.- β at $\beta = 7.0$ in the chaotic regime.

infinite time average to Eq. 2.40 [141],

$$\begin{aligned}
0 &= \left\langle \frac{dJ_{\text{ac}}}{dt} \right\rangle + \left\langle \int_0^1 \zeta p^2 dx + p_f \dot{q} \right\rangle \\
&= \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T \frac{dJ_{\text{ac}}}{dt} dt + \left\langle \int_0^1 \zeta p^2 dx \right\rangle - \langle p_f \dot{q} \rangle \\
&= \lim_{T \rightarrow \infty} \frac{J_{\text{ac}}(T) - J_{\text{ac}}(0)}{T} + \left\langle \int_0^1 \zeta p^2 dx \right\rangle - \langle p_f \dot{q} \rangle. \tag{2.43}
\end{aligned}$$

Considering that the acoustic energy is a bounded quantity on a strange attractor, the first term of the above equation is 0. Hence, Eq. 2.43 physically means that the damping mechanism exactly balances the acoustic source at regime, i.e.,

$$\langle J_{\text{ray}} \rangle = \langle p_f \dot{q} \rangle = \left\langle \int_0^1 \zeta p^2 dx \right\rangle. \tag{2.44}$$

Thus, the time-averaged Rayleigh index can be expressed either from the heat-source contribution or the dissipation term. From a computational point of view, the calculation of the sensitivity of $\langle p_f \dot{q} \rangle$ is difficult because the chaotic modulation, which is imposed exactly at $x = x_f$, makes $\langle p_f \dot{q} \rangle$ erratic. To overcome this computational problem, we recommend using $\left\langle \int_0^1 \zeta p^2 dx \right\rangle$ (bearing in mind the equality Eq. 2.44), which numerically behaves regularly because it is an integral quantity [141]. Using

the Galerkin modal decomposition of the pressure, this integral becomes Eq. 2.39, which was used earlier to define the Rayleigh index. In Figure 2-3, on the right hand side, we plot the ergodic average of the Rayleigh index, $\langle J_{\text{ray}} \rangle$ as a function of heat release β . The behavior of $\langle J_{\text{ray}} \rangle$ is consistent with that of $\langle J_{\text{ac}} \rangle$, with irregular values in the chaotic regime ($6.4 \leq \beta \leq 7.3$), and a sharp increase in the chaotic-to-periodic transition. Note that the cost functional $\langle J_{\text{ray}} \rangle$ is not directly proportional to the norm of the state, unlike the acoustic energy. The sensitivity and optimization framework we propose can tackle general cost functionals.

2.4.4 Lyapunov vectors in the chaotic regime

We treat the numerical solution of the system of ODEs as the map f between consecutive timesteps. The parameters \mathcal{S} , as per our notation in section 2.3, is set to $[\beta, \tau]^T$. Corresponding to the LEs (shown in Figure 2-3), we also compute the tangent Covariant Lyapunov Vectors (CLVs) using Ginelli *et al.*'s algorithm [113]. The CLVs are tangent/adjoint vectors whose asymptotic exponential growth or decay rates are exactly equal to the LEs; further, they are covariant in the sense that a homogeneous tangent/adjoint solution (Eq. 2.4 and Eq. 2.10) starting from a CLV always lies in the span of the same CLV field. The reader is referred to [159] for the properties of CLVs. In this chapter, the CLVs are normalized. The unstable (stable) tangent CLVs form a basis of (non-orthonormal) unit vectors for E^u (E^s); the span of the unstable (stable) adjoint CLVs is $(E^s)^\perp$ ($(E^u)^\perp$). Since the center subspace is one-dimensional, the normalized time-derivative $\mathcal{F}/\|\mathcal{F}\|$ is the center tangent CLV field. The Ginelli algorithm can also compute the adjoint CLVs. For this, we use the Jacobian transpose trajectory, in place of the Jacobian trajectory used to compute the tangent CLVs, and we also reverse time (i.e, the forward/backward phase of Ginelli's algorithm is carried out backward/forward). We compute the tangent and adjoint CLVs at $\beta = 7$. In Figure 2-5, we show the angles between each pair among the first 6 tangent CLVs, and each pair of adjoint CLVs, on the right. The angles are averaged over 250 time units. In a hyperbolic system, the angles between every pair of CLVs (corresponding to different LEs) are uniformly (in phase space) bounded away from 0. Although not




Figure 2-5: Ergodic average of the angles between the first 6 different adjoint CLVs (right) and tangent CLVs (left). The averaging window was set at 250 time units. The colorbar shows the angle in degrees.

a rigorous test for hyperbolicity, the results of Figure 2-5 indicate that, at least on average, the first 6 CLVs, both tangent and adjoint, do not show tangencies. Over the time window of calculation, the minimum angle observed between any dissimilar pair was about 4 degrees. This indicates that the system is likely uniformly hyperbolic. Biorthogonality is numerically verified in Figure 2-6. Except along the diagonals, which contain the mean angles between a tangent and adjoint CLV corresponding to the same LE, the mean angles are all about 90 degrees, as expected.

2.5 Suppression of a nonlinear oscillation by gradient-based optimization

The thermoacoustic model under investigation displays chaotic behavior in the region $6.4 \leq \beta \leq 7.3$. As illustrated in Figure 2-1, conventional methods to compute the sensitivities of the long-time behavior of this model, in this chaotic regime, fail to produce meaningful sensitivities. In this section, we use the shadowing algorithm presented in section 2.3 to enable the computation of these sensitivities. Our goal is to illustrate the potential of the algorithm for practical sensitivity-based optimization and parameter estimation in the regime of chaotic acoustics.

Given the small number of parameters and objective functions, it is possible to




Figure 2-6: Angles between pairs of adjoint and tangent CLVs when averaged over 250 time units. As expected, each tangent CLV is orthogonal to every adjoint CLV except those with the same LE.

compare the sensitivities computed through the algorithm with the slopes obtained from the bifurcation diagrams in Figures 2-2 and 2-3 (right). These comparisons validate the results of our algorithm, which is one of the goals in this section. We demonstrate the usefulness of the computed sensitivities by using them in a gradient descent algorithm to minimize the ergodic average of the acoustic energy. This simple optimization procedure can be used for heat release parameter selection. The NILSS algorithm and its discrete AD variant presented here thus introduce sensitivity-based optimization and parameter estimation to the chaotic regime, more generally in hyperbolic systems with constant time-delays, extending the work of Huhn and Magri [141].

2.5.1 Shadowing sensitivities of the acoustic energy and Rayleigh criterion

We apply the shadowing algorithm from section 2.3 to compute both tangent and adjoint shadowing sensitivities. We compute the tangent shadowing direction once to estimate $(d\langle J_{ac} \rangle / d\beta)$ and $(d\langle J_{ray} \rangle / d\beta)$. Similarly, we compute an adjoint shadowing direction once, to calculate both the sensitivities $d\langle J_{ac} \rangle / d\beta$ and $d\langle J_{ac} \rangle / d\tau$. Next we




Figure 2-7: Sensitivities of $\langle J_{ac} \rangle$ (red) and $\langle J_{ray} \rangle$ (blue) computed through tangent shadowing. The sensitivities are obtained by a cumulative average over sensitivities each of which is computed over 20 time units.

define the inputs to the shadowing algorithms. A primal orbit $\{u_n\}$ is a sequence of 30-dimensional solution vectors obtained by time-integrating the primal system (Eq. 2.33 and Eq. 2.36). The map f is the Tsitouras Runge-Kutta time-integrator that advances a solution state by one timestep. The timestep size is fixed at 0.01. For tangent shadowing, the input b_n is set to $x_n = (\partial f / \partial \beta)(u_{n-1}, \mathcal{S})$. In the adjoint shadowing algorithm, we arrange the input sequence $\{b_n\}$ so that $b_{n'+1}$ is set to $x_n^* = DJ_{ac}(u_n)$. (i.e., we pass x_n^* in time-reversed order to the adjoint shadowing algorithm). We use the AD package `Zygote.jl` [226] [225] to compute the sequence $\{b_n\}$ for tangent shadowing, through AD. The Jacobian matrix $A_n := Df(u_n)$ and its transpose, needed for the tangent and adjoint algorithms respectively, are computed by using finite difference. Instead, if using AD, each A_n must be computed row-by-row (since `Zygote.jl` does not support vector-valued outputs), leading to a much larger computation time compared to using finite difference. Moreover, the shadowing algorithms do not need the Jacobian to be computed exactly. In the AD version of tangent and adjoint shadowing, we do not need to compute A_n and b_n ; the tangent/adjoint perturbations needed in the n -loop (section 2.3.2) are directly computed using AD as shown in 2.10. The input to the AD version of both tangent and adjoint shadowing are the functions that perform primal time-integration and compute the objective functions, given the primal state. In order to ensure that u_0 is a point on the chaotic attractor, we evolve the system for a time of 10000 time units, starting from a random 30 dimensional vector.

We set $d_u = 2$ for both algorithms. Although we could theoretically have used




Figure 2-8: Sensitivities of $\langle J_{ac} \rangle$ with respect to τ (green, right) and $\langle J_{ac} \rangle$ (red, left), wrt β , computed through adjoint shadowing at $\beta = 6.9, \tau = 0.2$. The recorded sensitivities were obtained by a cumulative average over sensitivities each computed over a time length of 20 units.

$d_u = 1$ in the tangent algorithm, setting $d_u = 2$ leads to better approximations of the shadowing direction via the least squares problem (section 2.3.3). Each minimization problem is solved over a time duration of 20 time units, which is about 4 Lyapunov times. To obtain the sensitivity of a long-time average, a sample mean of these intermediate-time sensitivities is taken. A cumulative mean of the sensitivities converges as the length of time (number of samples) increases (2-7 and 2-8). In Figure 2-7, we show the sensitivities of the long-time averaged acoustic energy (left) and Rayleigh index (right), with respect to β , computed using sample averages of the tangent shadowing sensitivities. In Figure 2-8, we show the sensitivities of the long-time averaged acoustic energy with respect to β (left) and τ (right), computed using sample averages of the adjoint shadowing sensitivities. The mean up to a time of 200,000 (i.e., calculated using 10,000 shadowing sensitivities each over a time of 20 units) is shown as a solid line. The mean values shown in both plots in Figure 2-7 compare well against the corresponding slopes from the inset plots of Figures 2-2 and 2-4. The mean value of the sensitivity with respect to β , about -3.9, also agrees well, with the same sensitivity computed using the adjoint algorithm, around -3.5, which is shown in Figure 2-8 (left). Since all three quantities, the shadowing sensitivities from the two algorithms as well as a direct reading of the slope from the $\langle J_{ac} \rangle$ -vs- β plots, suffer from statistical noise due to a finite computation window, we do not expect exact agreement. Nevertheless, we note that both tangent and adjoint sensitivities are within 12% of the slope estimate of -4 obtained from by approximating $\langle J_{ac} \rangle$ -vs- β




Figure 2-9: The path of optimization starting at $\beta = 6.5$. The points are superimposed on the plot of $\langle J_{ac} \rangle$. In the gradient descent algorithm, the step size is taken to be 0.1.

as a line in Figure 2-2 (inset). To ensure its correctness, the program that implements the two shadowing algorithms is validated by computing sensitivities on the classical model of a chaotic ODE, the Lorenz'63 system, to verify the computed sensitivities against the values available for this system in the literature. For details on this validation and on the replication of Figures 2-8 and 2-7, see 2.8. The absolute values of the sensitivities $d\langle J_{ray} \rangle/d\beta$ from Figure 2-7 are consistent with our expectation from Figure 2-4 that the ergodic average $\langle J_{ray} \rangle$ decreases with β , around $\beta = 6.9$, but not as rapidly as $\langle J_{ac} \rangle$. From Figure 2-8 (right), we see that the ergodic average of the acoustic energy is highly sensitive (when compared to changes in β) to small perturbations in the time delay parameter τ . The shadowing algorithm shows convergence, when the same derivative is computationally prohibitive to obtain accurately with ensemble sensitivity calculations [96, 69].

2.5.2 Minimization of the acoustic energy using shadowing

We shall demonstrate an application of the sensitivity $d\langle J_{ac} \rangle/d\beta$ that we computed in the previous section, to the problem of optimal design by parameter selection [141].

We use the numerically computed bifurcation diagram (Figure 2-2) as a qualitative check of our parameter selection procedure. We aim to solve the following optimization problem,

$$\begin{aligned} & \underset{\beta}{\text{minimize}} && \langle J_{\text{ac}} \rangle(\beta) \\ & \text{subject to} && \text{Eq.2.36, Eq.2.33,} \end{aligned} \tag{2.45}$$

where $\langle J_{\text{ac}} \rangle$ is the ergodic average of the acoustic energy, which is approximated using a long-time average over a time duration of N steps. The parameter can be updated by a steepest descent method

$$\beta_{n+1} = \beta_n - \gamma d_{\beta} \langle J_{\text{ac}} \rangle |_{\beta=\beta_n}, \tag{2.46}$$

A relaxation factor of $\gamma = 0.1$ is used to enable stable and accurate numerical convergence. At each β , the shadowing sensitivity is computed over 1000 time units, which is a sample average of 50 sensitivities collected from runs of the shadowing algorithm each over a time of 20 units. The algorithm stops when the condition

$$\langle J_{\text{ac}}(\beta_n) \rangle < \epsilon \cdot \langle J_{\text{ac}}(\beta_0) \rangle \tag{2.47}$$

is met, where $\epsilon = 1\%$. This condition physically signifies that the optimization is successful when the system vibrates around the fixed point. In Figure 2-9, we show the optimization path (blue points) taken by this procedure, starting from $\beta = 6.5$. We show the points $(\beta, \langle J_{\text{ac}} \rangle(\beta))$ numbered in the order in which they are encountered in the optimization procedure. As shown in the figure, the path leads to a reduction in $\langle J_{\text{ac}} \rangle$ by exiting the chaotic region into the periodic regions for larger of β , and eventually into the periodic region at smaller values of β . We remark that, in this case, direct evaluation of the bifurcation diagram is possible, owing to the relatively low dimension of the system and the parameter space. This bifurcation diagram itself suggests optimal paths for acoustic energy reduction. However, in a more general setting, the dimension of the system, the objective function and parameter spaces may be such that bifurcation diagrams, at the resolution of parameters required to

compute accurate gradients, are computationally infeasible. The demonstration in this section indicates that the shadowing sensitivities, which can be computed at a smaller cost relative to the bifurcation diagram computation, can be used instead, for optimization problems. For illustration purposes, we have chosen a simple gradient-descent with a fixed relaxation factor. The numerical results in Figure 2-9, however, show that the step size plays an important role in the rapid transition from large β in the periodic region (around 7.4) to the periodic region for smaller values (around 4.0). We defer to a future work the effect of the step size, which is beyond the scope of this chapter.

2.6 Data assimilation with discrete shadowing

A common problem whenever we have incomplete, and often noisy, observations together with a model of a physical system, is to estimate a model trajectory that reproduces the observations. In this section, we explain that the variational formulation of this problem is an application of the shadowing algorithm discussed in section 2.3. Then, we apply the shadowing algorithm to the time-delayed model to illustrate its potential for data assimilation in chaotic solutions.

In a data assimilation problem, we are given external measurements of an observable g at a sequence of times, denoted $g_{n_k}^{\text{obs}}$, $1 \leq k \leq N$. In the data assimilation method 4DVar [259], an initial state u_0 is sought so that the model trajectory at the observation times, $f_{n_k}(u_0) = u_{n_k}$ produces a sequence $g_{n_k} := g(u_{n_k})$ that closely matches the observations $g_{n_k}^{\text{obs}}$. If a reliable guess for the initial state, known as *background* u_0^{bg} , is available, we desire our predicted initial state to be close to the background. The optimal initial state, known as *analysis*, is obtained by minimizing the following cost functional

$$\langle J \rangle(u_0) = \frac{1}{2}(u_0 - u_0^{\text{bg}})^T B^{-1}(u_0 - u_0^{\text{bg}}) + \frac{1}{2} \sum_{k=1}^N (g_{n_k}^{\text{obs}} - g_{n_k})^T L^{-1}(g_{n_k}^{\text{obs}} - g_{n_k}). \quad (2.48)$$

The first term in the cost functional in Eq. 2.48 corresponds to the misfit between

the predicted initial state and the background, where B is a $d \times d$ background error covariance matrix. The second term is the misfit between the external observations and values of the observable generated by the model, using the predicted state as the initial condition. Given an observable space of dimension l , the $l \times l$ matrix L is the observation error covariance matrix. Without loss of generality, take B and L to be identity matrices (in $\mathbb{R}^{d \times d}$ and $\mathbb{R}^{l \times l}$ respectively), which results in the following cost functional:

$$\langle J \rangle(u_0) = \frac{1}{2} \left\| u_0 - u_0^{\text{bg}} \right\|^2 + \frac{1}{2} \sum_{k=1}^N (g_{n_k}^{\text{obs}} - g_{n_k})^2. \quad (2.49)$$

The cost functional is minimized using a standard optimization procedure, and an analysis state u_0^* is obtained. The reader is referred to [260] for data assimilation of nonchaotic states with a similar thermoacoustic model considered in the present chapter. When the model is chaotic, the standard optimization procedure fails since the gradient of the functional with respect to u_0 grows exponentially with the time duration of available observations (Figure 2-1). Thus, assimilation for a time window longer than the typically short Lyapunov time, cannot be achieved using standard optimization methods.

However, several effective strategies have been proposed, particularly in the field of numerical weather prediction, wherein chaotic models are widely used, that are successful over long assimilation windows. The most popular of these include 4DVar-AUS, in which the *analysis increment* – the discrepancy added to the state during optimization – is restricted, to the nonstable ($E^u \oplus E^c$) subspace, at every timestep [262]. Another method is projected shadowing-based data assimilation in which the cost functional is minimized using Newton’s method in each step of which the analysis trajectory, as a whole, is updated. Additionally, the updates to the analysis trajectory during each Newton iteration is carried out only on the nonstable subspace, and this leads to the economy of the method. The updates to the analysis trajectory on the stable subspace is treated using a different method, known as synchronization [83]. In this work, we present an alternative approach that uses the shadowing algorithm

described in section 2.3. The goal of this approach is to compute a model orbit that shadows the pseudo-orbit pertaining to the observations. Our method thus joins the class of shadowing-based data assimilation methods offering an alternative formulation that indirectly computes the shadowing orbit through NILSS.

2.6.1 Tangent NILSS for state estimation with full-state observations

We introduced NILSS [206] as a method to differentiate long-time averages with respect to parameters in a chaotic system. How is the method relevant to the problem of state and parameter estimation? The answer to this question lies in the fact that a shadowing trajectory can be obtained as a byproduct of the NILSS method, and shadowing trajectories can be used for state estimation. First, we recognize that, in NILSS, the derivative of the long-time average is computed along a shadowing trajectory at a perturbed parameter. Secondly, we relate the shadowing perturbation sequence, $\{v_n^{\text{sh}}\}$, that is computed by tangent NILSS, to the solution of the state and parameter estimation problem.

Shadowing-based interpretation of NILSS

NILSS and its adjoint versions use the shadowing lemma (see e.g. 18.1.2 of [154]) by considering the reference trajectory $\{u_n\}$ as a pseudo-orbit of $f(\cdot, s + \epsilon)$. According to the shadowing lemma, there is a unique orbit of $f(\cdot, s + \epsilon)$ called the shadowing orbit, that is close to the given reference orbit of $f(\cdot, s)$, u_n . The tangent *shadowing perturbation* is the sequence of tangent vectors, along this shadowing orbit, that expresses the discrepancy between the shadowing orbit and the pseudo-orbit $u_n, n \in \mathbb{Z}^+$, in the limit $\epsilon \rightarrow 0$. Hence it remains bounded for all time. A close approximation of the shadowing perturbation, v^{sh} , is obtained by solving the following least squares

problem ([271], Theorem LSS)

$$\begin{aligned}
v^{\text{sh}} &= \operatorname{argmin}_{w \in \mathcal{V}} \sum_{n=0}^{N-1} \|w_n\|^2 \\
\text{s.t. } w_{n+1} &= D_u f(u_n) w_n + x_{n+1}.
\end{aligned} \tag{2.50}$$

In the above optimization problem, the search space \mathcal{V} is restricted to a subset of $(\mathbb{R}^d)^N$ in which each w_n can be expressed as $w_n = v_n + Q_n a_n$, i) for some $a_n \in \mathbb{R}^{d_u}$, ii) where v_n is the conventional tangent solution, i.e., solution of Eq. 2.3 with zero initial condition. This leads the NILSS algorithm [206] to being more efficient than if the search space was set to $(\mathbb{R}^d)^N$ [206].

2.6.2 Converting the state estimation problem to a parameter estimation problem

The cost functional in state estimation consists of two parts: the background error and the observation error. We may assume that the background orbit is a pseudo-orbit of $f(\cdot, s + \epsilon)$, for some ϵ around zero. Then, the shadowing problem (Eq. 2.50) minimizes the first part of the state estimation cost functional, which is the background error. In order to minimize the second part, the observation error, we use the observation error as the objective function in the shadowing algorithm. Through the parameter optimization procedure described in section 2.5, we find an optimal parameter that minimizes the observation error. This amounts to finding an ϵ such that the error between an observation orbit and a true orbit of $f(\cdot, \epsilon)$ is minimized. Then, we can refine the background trajectory, using the shadowing perturbation (computed by the shadowing algorithm). The parameter optimization is then repeated with the refined trajectory as the new background. By repeating this combined state-parameter optimization procedure, we minimize both parts of the state estimation cost functional separately. At the end of this procedure, the analysis state is obtained by iteratively refining the background state using the shadowing perturbations at different parameters. We outline the state-parameter optimization procedure assuming

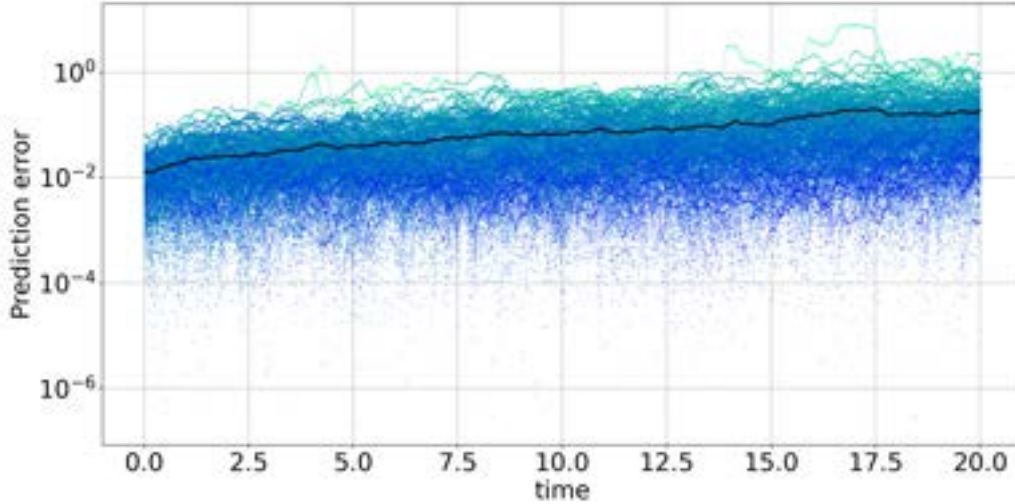


Figure 2-10: Relative error in the instantaneous acoustic energy between its predicted and observed values in the chaotic regime ($\beta = 7$). The maximum error increases from blue to green colors. The mean error (over the assimilation time window) is shown in black.

we have observations at every timestep. The following steps are repeated, until the cost functional is less than a specified tolerance. Initially, the reference orbit u_n for shadowing is set to the background orbit. The algorithm is the following:

1. Run tangent shadowing (section 2.3) with an objective function

$$\langle J \rangle(s) := \frac{1}{N} \sum_{n=0}^{N-1} (g_n^{\text{obs}} - g_n)^2, \quad (2.51)$$

to obtain $d_s \langle J \rangle$, and the sequence $\{v_n^{\text{sh}}\}$.

2. Update the parameter as: $s \leftarrow s + \delta s$, where $\delta s = \gamma d_s \langle J \rangle$.
3. Update the initial condition for the next iteration as $u_0 \leftarrow u_0 + \delta s v_0^{\text{sh}}$. Go to step 1.

In practice, the time 0 corresponding to the start of the assimilation window, must be postponed roughly by $1/\lambda_1$ (Lyapunov time), to allow a spin-up time for v_0^{sh} to be accurately computed. The relaxation factor γ is assumed to be a fixed constant, as in the parameter optimization procedure in section 2.5.

2.6.3 Numerical results

We present state estimation results on the time-delayed model, computed using the above algorithm. The algorithm is validated on the Lorenz'63 model, as shown in Figure 2-12 in 2.9. A total of 180 experiments are performed on the time-delayed model, each with a different background trajectory of length 20 time units. The parameter s that is updated is set to β , with the reference value of 7.0. In each experiment, a background state is generated by perturbing each component of a reference state by a Gaussian random variable of variance 0.1. The original trajectory, started from the reference state, is used to generate the observations (of the acoustic energy) J_{ac}^{obs} at every timestep. The mean squared observational error is the objective function for tangent NILSS,

$$\langle J \rangle := \frac{1}{N} \sum_{n=0}^{N-1} |J_{ac_n}^{\text{obs}} - J_{ac_n}|^2.$$

The step size for gradient descent (section 2.5) is $\gamma = 0.1$. We show in Figure 2-10 the relative errors in the acoustic energy along the analysis orbit, which is the result of the algorithm after 200 gradient descent steps. The relative error is defined as $|J_{ac_n}^{\text{obs}} - J_{ac_n}|/J_{ac_n}^{\text{obs}}$. Each colored line indicates a single experiment, with a total of 180 numerical experiments performed with different background states. The color of the line indicates the maximum relative error observed in that experiment, during the assimilation window of 20 units; the maximum errors increase from blue to green. The sample mean of the relative errors across all the experiments is shown in black. Note that the assimilation time is 4 times longer than $1/\lambda_1$, the Lyapunov time. As shown in Figure 2-1, over 20 time units, we expect a small perturbation introduced in the initial condition, in almost any direction, to grow by 3 orders of magnitude. However, as the results in Figure 2-10 indicate, the relative error in J_{ac} has been restricted to within 10 % over this assimilation window, due to use of shadowing directions to iteratively refine the initial condition and the parameter, to match the observations. Improvements to this algorithm in order to reduce the errors, and further increase the predictability window will be studied in future work. One modification to the

suggested algorithm, toward this goal, is to incorporate the observation error at every timestep, into the perturbation x_n in the shadowing algorithm.

2.7 Conclusions

Naïve applications of linear perturbation methods such as tangent/adjoint/Automatic Differentiation (AD)/finite-difference, cannot compute the derivatives of long-time averages in chaotic systems with respect to specified inputs. A recent method, known as the Non-Intrusive Least Squares Shadowing [206], computes these derivatives by numerical construction of tangent/adjoint *shadowing* perturbations, which are infinitesimal perturbations that remain bounded for a long time duration. In this chapter, we introduce AD into the tangent and adjoint NILSS. This nontrivial combination of algorithms is an enabler for the application of shadowing to complex dynamical systems. We demonstrate shadowing by computing sensitivities on a chaotic time-delayed model, which is a reduced-order model of a gas turbine combustor. We compute tangent and adjoint shadowing sensitivities of the ensemble averages of the acoustic energy and Rayleigh index with respect to the design parameters that control the heat release rate and the time delay. Although the model is a reduced representation compared to CFD-based combustion models, it can be used to estimate the development of chaotic acoustic instabilities. First, we demonstrate an automatically-differentiated procedure for the minimization of the long time-averaged acoustic energy through heat-release parameter selection, which does not require a tangent solver [141]. Secondly, we construct a pseudo-orbit data assimilation scheme using the computed shadowing sensitivities in an optimization loop. We show that this scheme extends the predictability window by four Lyapunov times. Finally, we remark that the proposed algorithm, the shadowing-based optimization and data assimilation scheme are more generally applicable. The algorithms and the software developed (available at [58]) may be used for other hyperbolic chaotic models with/without time delay.

2.8 Validation of sensitivity computation on the Lorenz’63 model

We use the classical model of chaos, the Lorenz’63 system, for validation results in this chapter. The Lorenz’63 model is the following set of nonlinear ODEs that serves as a reduced-order model for fluid thermal convection between parallel plates maintained at a temperature difference [179]:

$$\frac{d}{dt} \begin{bmatrix} x \\ y \\ z \end{bmatrix} = \begin{bmatrix} 10(y - x) \\ x(s - z) - y \\ xy - (8/3)z \end{bmatrix}. \quad (2.52)$$

Here a phase point u is represented using 3 coordinates as $u \equiv [x, y, z]^T \in \mathbb{R}^3$ and the map $f(\cdot, s)$ advances a state u_n to u_{n+1} by timestepping the ODE system (Eq. 2.52). For the time integration, we use a Forward Euler scheme with a timestep of 0.005. The Lyapunov exponents of this system are about 0.9, 0, and -14.6. This is a partially hyperbolic system [14] that has been shown to possess an SRB measure [136]. The objective function for validating shadowing is chosen to be the z coordinate function. It is known that the ergodic average $\langle z \rangle$ as a function of s , can be approximated as a straight line with a slope of about 0.96, over a range of values around the standard $s = 28$ [167]. We use the discrete shadowing algorithms described in section 2.3 in both tangent and adjoint mode, to compute this derivative. Numerical results are shown in Figure 2-11, in which 100 sample sensitivities are shown each over a time duration of 15 units, starting from different points on the Lorenz’63 attractor. On the left, we see the computed tangent shadowing sensitivities with $d_2 = 2$, and the on the right, the adjoint shadowing sensitivities, also computed with $d_u = 2$. From the numerical results, we see that the sample means of both sensitivities are within 10% of the reference value of 0.96, thus validating both tangent and adjoint shadowing sensitivity codes. These tests can be run from `tests/test_lorenz63.jl` at [58]. Interestingly, adjoint shadowing appears to be better suited for this objective




Figure 2-11: Sensitivities $d\langle z \rangle / ds$ computed for the Lorenz'63 model using the tangent algorithm (left) and the adjoint algorithm (right). Different trajectories of length 3000 or 15 timeunits are used to compute the sensitivities, shown as an errorbar of length one standard deviation.




Figure 2-12: Relative error in the predicted state z as a function of time for the Lorenz'63 system. The maximum prediction error increases from light to dark colors. The mean error across all experiments is shown in black.

function-parameter pair since the variance of tangent sensitivities is 10 times larger than the adjoint.

2.9 Validation of data assimilation scheme on the Lorenz'63 model

We apply the data assimilation scheme described in section 2.6 on the Lorenz'63 model. We perform a series of 100 numerical experiments each with a different back-

ground. We fix a reference trajectory of length 5000 steps (10 time units), and generate a background by perturbing the reference initial condition in the \hat{z} direction by an additive Gaussian noise with a variance of 0.1. The reference trajectory is used to generate the observation trajectory $J_n^{\text{obs}} := z_n^{\text{obs}}$. Tangent NILSS is run with an objective function,

$$\langle J \rangle(s) := \frac{1}{N} \sum_{n=0}^{N-1} (z_n^{\text{obs}} - z_n)^2, \quad (2.53)$$

where z_n is the z -coordinate of a trajectory u_n . Using the derivative $(d\langle J \rangle/ds)(s)$ computed by shadowing, the parameter s is updated in a gradient descent algorithm with a constant step size of $\gamma = 0.1$. At the beginning of gradient descent, the trajectory is set to the background trajectory. At each gradient descent step, it is iteratively refined to a pseudo-orbit, as described in Step 3 in section 2.6. The mean error is calculated by averaging the optimal error across the 100 experiments. The optimal errors, defined as relative observation errors after 200 gradient descent steps, across experiments are as shown as a function of time in Figure 2-12. The relative observation error at time n is $|(z_n^{\text{obs}} - z_n)|/z_n^{\text{obs}}$. The colors of the lines in Figure 2-12 are according to the maximum over n of the relative observation error. As mentioned before, the only positive Lyapunov exponent of this system is known to be about 0.9. The results indicate a predictability within 10% of the observation, on average, even up to 10 time units, which is about $10 \times (1/\lambda^1)$. With this validation on the Lorenz'63 model, we apply the same method for data assimilation in the Rijke tube model, as described in section 2.6.

2.10 AD shadowing and code for replication/extension

As we describe in the main text, the AD version of shadowing is enabled by the introduction of AD to replace the tangent/adjoint solvers in tangent/adjoint shadowing. In this section, we give a pseudocode for this modification. The source code in Julia for generating the numerical results in this chapter is available at [58]. This section also

briefly describes how this code may be used with a different chaotic ODE/map. We must mention that an alternative AD-version of the NILSS algorithm is also available at [68], in which the `OpenAD` AD software package [266, 6, 138], written in Fortran, is used. This latter code assumes that the primal solver is available from the user as a binary file.

In the AD version of tangent/adjoint shadowing, as mentioned before, we replace the needed tangent/adjoint solvers with AD. In order to do this, we define a function, say \mathbf{f} , whose return value will be differentiated in forward/reverse-mode, and we specify some input variable(s) to differentiate with respect to. The AD software returns the gradient of the return value with respect to the input which is a function of the arguments passed to \mathbf{f} . Let \mathbf{f} be a function with the arguments x, y . The scalar or vector return value of this function is written as $\mathbf{f}(x, y)$. For instance, using the `Zygote.jl` Julia package, one would use the following syntax to obtain the gradient of $\mathbf{f}(x, y)$ with respect to x : `Zygote.gradient(x-> f(x, y), x)`. If using the `OpenAD` Fortran package instead, we declare the input x and the output of \mathbf{f} to be *active* variables. If `ret_f` is the variable that stores the return value of \mathbf{f} , its derivative with respect to x is stored in the variable `ret_f%d`. In the AD version of tangent/adjoint shadowing, we only propose to modify the n -loop in section 2.3.2 by introducing AD to replace the tangent/adjoint solvers; the rest of the algorithm is retained as described in section 2.3. The inputs to the AD shadowing algorithm are the sequences $\{u_n\}$ and $\{J_n\}$. We now give the modified n -loop in which we define the needed function, and input variables, for each AD invocation.

1. Obtain q_n^i from q_{n-1}^i by applying AD. In particular,
 - for tangent shadowing, we differentiate the function $f(u_{n-1} + \epsilon q_{n-1}^i, s)$. The input variable is ϵ and the output is the return value of the function. The derivative obtained from forward-mode AD is q_n^i , and this must be carried out for $1 \leq i \leq d_u$;
 - for adjoint shadowing, we differentiate the value $f(u_{n'+1}, s) \cdot q_{n-1}^i$, where $n' = N + 1 - n$. The input variable is $u_{n'+1}$. The derivative obtained

from reverse-mode AD is q_n^i , and this must be carried out for $1 \leq i \leq d_u$.

2. Let Q_n be the matrix with the columns q_n^i , $1 \leq i \leq d_u$. QR-factorize Q_n and set Q_n to the obtained “Q”. Let the “R” from QR factorization be stored as R_n . Thus, each q_n^i , $1 \leq i \leq d_u$ is now a unit vector.
3. Obtain v_n from v_{n-1} by applying AD. In particular,
 - for tangent shadowing, we differentiate the function $f(u_{n-1} + \epsilon v_{n-1}, s + \epsilon)$. The input variable is ϵ . The derivative obtained from forward-mode AD is v_n ;
 - for adjoint shadowing, we differentiate the value $v_{n-1} \cdot f(u_{n'+1}, s) + (1/N)J_{n'+1}$. The input variable is $u_{n'+1}$. The derivative obtained from reverse-mode AD is v_n .
4. Set $\pi_n := Q_n^T v_n$, which is a d_u -length orthogonal projection row-vector of v_n along q_n^i .
5. Normalize v_n by projecting out the components along q_n^i , that is, set $v_n \rightarrow v_n - \pi_{qv} Q_n$.
6. Go to step 1 with $n \rightarrow n + 1$, and repeat until $n = N$.

In the code at [58], the Julia script `tests/test_lorenz63.jl` and `tests/test_rijke.jl` compute respectively, for the Lorenz’63 model (Eq. 2.52) and the Rijke tube model (Eq. 2.33), the tangent and adjoint shadowing sensitivities. In order to use the code at [58] with a different model, we can use one of these test files as a template. The Julia file describing the model equations, which must be included in the test file, can follow the existing examples in `examples/lorenz63.jl` or `examples/rijke.jl`. The optimization and data assimilation routines described in section 2.5 and 2.6 respectively, can also be applied to a new model. To do this, we include the file containing the model equations into the utilities i) `utils/optimize.jl` for parameter optimization and ii) `utils/rijke_tangent_state_estimation.jl` for the state estimation algorithm.

Chapter 3

On the probability of finding a nonphysical solution through shadowing

This chapter proves that shadowing solutions can be almost surely nonphysical. This finding invalidates the argument that small perturbations in a chaotic system can only have a small impact on its statistical behavior. This theoretical finding has implications for many applications in which chaotic mechanics plays an important role. It suggests, for example, that we can control the climate through subtle perturbations. It also suggests that numerical simulations of chaotic dynamics, such as turbulent flows, may fail to predict the true long-term or statistical behavior. This chapter has been published at [66], and a preprint of this chapter can be found at [61].

3.1 Introduction

In this chapter, we show that a large mismatch is possible in the statistical or long-time averages measured in a pair of chaotic solutions to slightly different models. We construct numerical and analytical examples that illustrate this feature of chaotic dynamics. The analysis of the long-term behavior of *shadowing* solutions in the majority of the chaotic models constructed in this chapter reveals that shadowing solutions are

almost surely *nonphysical*: they do not represent the behavior of a physically observed solution starting almost everywhere. This also leads us to the conclusion that shadowing-based sensitivities [273] of long-time averages to parameter changes can be incorrect. Thus, the main conjecture of this chapter urges us to reconsider the use of chaotic numerical simulations, and statistical sensitivity analysis on these simulations. These implications are particularly important for climate studies, where chaotic models are widely used to make long-term predictions. In particular, three implications of the evidence in this chapter may be stated, the third of which harnesses the statistical inconsistency between slightly different models:

- we cannot trust the long-time averages generated by numerical simulations of chaotic processes, even when we eliminate model uncertainties stemming from imperfect modeling of the underlying physics, and statistical noise due to finite-time averaging;
- we cannot trust the sensitivities of long-time averages or statistical averages to small parameter changes computed by using shadowing-based methods for sensitivity computation;
- it is possible for two models separated by tiny changes in parameter values to have drastically different statistical behavior. This means that the Lorenzian butterfly can significantly alter the climate (statistical/long-term average) of Texas, and not just its weather (short-term events).

The chapter is organized as follows. In section 3.2, we introduce the dynamical systems concept of shadowing and its applications in the context of chaotic numerical simulations. In section 3.3, we define the notion of physicality and give examples to illustrate the different ways in which to obtain nonphysical solutions. Section 3.4 is dedicated to several examples of perturbed tent maps that have nonphysical shadowing solutions. In the discussion in section 3.5, we illustrate with an example that a small parameter perturbation can change the statistics significantly. The supplementary material, section 3.6, contains the numerical procedure to generate the probability distributions of the shadowing solutions.

3.2 Shadowing and applications

Our main analysis tool in advancing our conjecture is *shadowing*. Shadowing and other dynamical systems concepts used in this chapter are introduced in detail, where they appear, so as to make our presentation accessible to the large audience of computational scientists across different disciplines. Shadowing refers to the relationship between a pair of solutions to slightly different governing equations. The difference between the governing equations can be due to parameter perturbation or numerical error. The solution to one governing equation is said to shadow a solution to a second, slightly different equation if the first solution stays close to the second solution for some amount of time.

When the slightly different governing equations are uniformly hyperbolic, it has been long known that infinitely long shadowing solutions exist [10][46]. Numerical experiments in non-hyperbolic systems have shown that shadowing solutions may still exist [119][127] for a finite amount of time.

The existence of shadowing trajectories underlies many applications. Numerical simulations of turbulent flows have been widely used to study its statistical behavior. It is argued, based on shadowing, that such numerical simulations of chaotic dynamical systems can be useful, despite the *butterfly effect* that causes infinitesimal errors in the state to grow exponentially in time. Because of numerical and modeling errors, there is invariably a small difference between the true governing physics and the equations solved on a computer. By nature of chaos, this difference increases exponentially, as the system is evolved forward in time. Thus, the fidelity of numerical solutions of chaotic systems, such as turbulent flows, to the true physics is called into question. The existence of shadowing solutions is used to argue for the usefulness of such numerical solutions.

When certain conditions for shadowing theorems are met, the numerical solution would be an approximation to a “true” solution that satisfies the real governing physics.

Shadowing is also used in sensitivity analysis of chaotic dynamical systems. In

particular, it is used in computing how long-time-averages in a chaotic system respond to small perturbations in the governing equation [273][206][165]. In this application, the derivative of the long-time-averages is computed using a solution to the perturbed governing equation that shadows a solution to the unperturbed equation. The Least Squares Shadowing [273][206] method uses this concept.

There is an implicit assumption in both of these applications. The assumption is that the shadowing trajectory is a *physical* trajectory, a trajectory on which the long-time-average of a quantity is equal to the ensemble average. Not all solutions satisfying the physical governing equation are considered physical in this sense. In high-Reynolds number fluid flows, for example, a steady-state, laminar flow solution may satisfy the Navier-Stokes equation. But such a solution would never be observed in reality because it is unstable, and any small perturbation would trip it into turbulence. Unstable steady-state solutions are not the only nonphysical solution. Many chaotic dynamical systems have infinitely many periodic solutions that are, similar to their steady counterpart, unstable. These trajectories can have a probability distribution that is remarkably different from that of a typically observed solution of the governing equation. It is also possible to effect significant change in the statistics of the true governing dynamics by introducing a minor parameter perturbation. Thus, although trajectories of governing equations at slightly different parameters may shadow each other, the long-time-averages that are typically observed at one parameter may not be close to that of the other. As we will see in the next section, even solutions that look “chaotic”, i.e., unsteady and aperiodic, may not be physical.

3.3 What are physical and nonphysical solutions?

Intuitively, we call a solution to a governing equation *physical* if it represents what one would observe in a physical experiment. In particular, the statistics measured from a physical solution match the statistics observed in an experiment. Not all solutions are physical. A laminar flow solution, despite satisfying the governing equation, does not produce the turbulent statistics one would observe in a high-Reynolds number

experiment. Such solutions are thus called *nonphysical*.

In this section, we first describe what distinguishes, mathematically, a physical and a nonphysical solution. We also explain why it is theoretically unlikely to observe a nonphysical solution in an experiment. We then give a few examples of these almost-never-observed nonphysical solutions.

3.3.1 Physical solutions

What we have intuitively defined as physical solutions to a governing equation satisfy the following two criteria:

1. Time-averaged quantities converge in the limit of infinite averaging time. Consider $u(t)$ as the solution to a chaotic governing equation, then for a regular observable of interest $J(u(t))$,

$$\lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T J(u(t)) dt \quad (3.1)$$

exists.

2. For almost any small perturbation to the initial condition $u(0) \rightarrow u'(0)$, the perturbed solution $u'(t)$ should have the same statistics, i.e.,

$$\lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T J(u(t)) dt = \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T J(u'(t)) dt \quad (3.2)$$

The left column of Figures 3-1-3-3 illustrates a physical solution of the Lorenz system,

$$\frac{du}{dt} = \frac{d}{dt} \begin{pmatrix} u_1 \\ u_2 \\ u_3 \end{pmatrix} = \begin{pmatrix} \sigma(u_2 - u_1) \\ u_1(\rho - u_3) - u_2 \\ u_1 u_2 - \beta u_3 \end{pmatrix} \quad (3.3)$$

where $\sigma = 10$, $\rho = 28$, and $\beta = 8/3$. The solution starts at the initial condition $u(0) = (0.01, 0.01, 28)$. After a long time evolution, we observe that the solution has visited a large portion of the $u_1 - u_3$ plane with a varying, but well-defined,

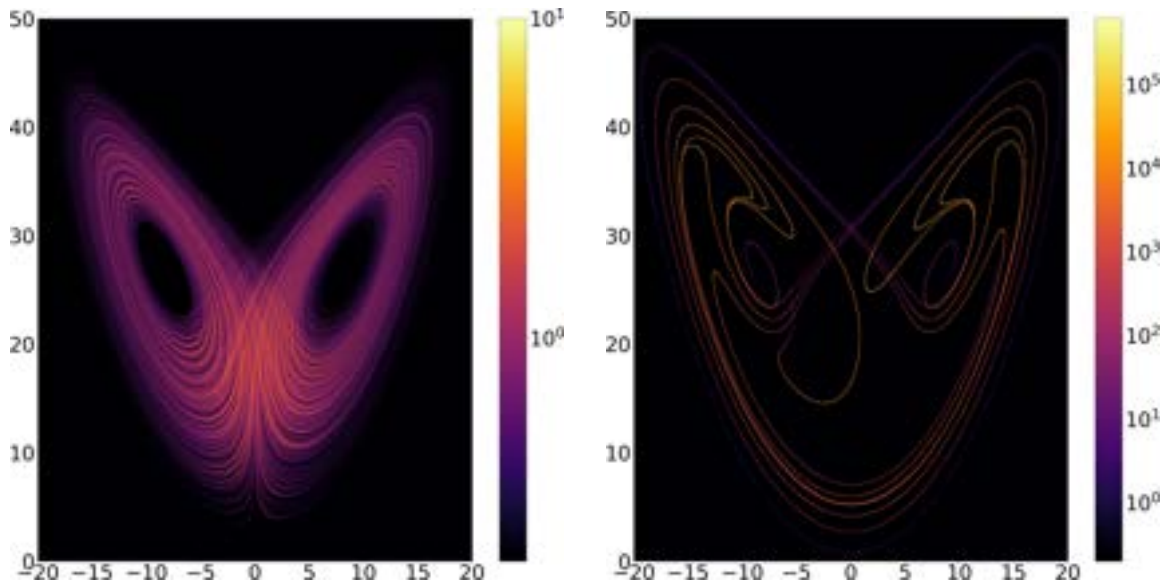


Figure 3-1: L: An ensemble of initial conditions distributed uniformly in the box $u_1 \in [0, 1], u_2 \in [0, 1], u_3 \in [28, 39]$ after 10 time units of evolution, shown on the $u_1 - u_3$ plane. R: distribution of a trajectory of 1000 time units in length. In both plots, the color represents the number of samples in a 2048×2048 uniform grid. The trajectory is sampled every 0.001 time units.




Figure 3-2: L: distribution of the same ensemble as in Figure 3-1 after another 5 time units of evolution. R: distribution of the same trajectory as in Figure 3-1 evolved for 10,000 time units.




Figure 3-3: L: distribution of the same ensemble as in Figures 3-1 and 3-2 after a total of 50 time units of evolution. R: distribution of the same trajectory as in Figures 3-1 and 3-2 evolved for 100,000 time units.

frequency. We can use a probability distribution, μ , to quantify how frequently a physical solution visits portions of the phase space. Specifically, for a subset A of the phase space, $\mu(A)$ measures the fraction of time a very long physical solution spends inside the subset A . With this distribution defined, the infinite-time average of any quantity $J(u)$ can be represented as an average of J over the entire phase space U , weighted by this statistical distribution μ . Mathematically,

$$\frac{1}{T} \int_0^T J(u(t)) dt \xrightarrow{T \rightarrow \infty} \int_U J(u) d\mu(u) \quad (3.4)$$

Remarkably, the distribution μ not only characterizes the history of a long physical solution, but also describes the settled state of an ensemble of solutions. This is illustrated on the right column of Figures 3-1–3-3. We generate these plots by starting from an ensemble of about one billion initial conditions, randomly and uniformly spaced in a small three-dimensional box. All these billion solutions are evolved by solving Eq. 3.3 for 10, 15, and 50 time units, to obtain the plots. We observe that, as time evolves, the ensemble spreads over an increasingly larger portion of the $u_1 - u_3$ plane. After a long time, the ensemble settles into a time-invariant *attractor* that

contains the *unstable manifold*, which appears as filaments forming the attractor. The probability distribution of the ensemble on the attractor becomes identical to the distribution of a single, very long physical solution, which is also contained in the attractor (Figure 3-3).

This remarkable agreement has been thoroughly studied under the subject of ergodic theory. Under surprisingly weak conditions, a solution starting from **almost** any initial condition, chosen randomly from a set enclosing the attractor, is a physical solution [281]. Meanwhile, an ensemble of trajectories starting from any distribution with a finite density also evolves towards the same final distribution, μ . Due to expansion of a volume of solutions tangent to the attractor filaments, a finite density under long-time evolution becomes absolutely continuous on the unstable manifold, i.e., the likelihood of a trajectory visiting any set not intersecting the attractor filaments is zero. The stationary distribution achieved on long-time evolution of the ensemble that has the absolute continuity property is called the Sinai-Ruelle-Bowen (SRB) measure. The absolute continuity property is sufficient to ensure that the SRB measure is the same μ . That is, the SRB measure is physically observed in the sense that physical solutions produce long-time-averages which are expectations with respect to the SRB measure.

Note that almost any, not any, initial condition leads to a physical solution. A set of special initial conditions contained in a neighborhood of the attractor may exist starting from which physical solutions are not generated. These initial conditions do not produce the same statistics as the physical solutions. This special set of initial conditions is Lebesgue measure zero – one has zero chance of finding such an initial condition by randomly sampling. Nonphysical solutions thus take an effort to find. Nevertheless, they turn out to be important when discussing shadowing, the topic of this chapter. We first introduce a type of obviously nonphysical solutions in the next subsection, before discussing a less obvious type in section 3.3.4.




Figure 3-4: Periodic solutions of the Lorenz equation, overlaid on top of its SRB distribution. The solid line, dashed line, and dotted line represent three distinct periodic solutions.

3.3.2 Nonphysical solutions type I: Periodic solutions

A periodic solution is nonphysical because it does not visit as much of the phase space as a physical solution does. Figure 3-4 shows a few periodic solutions of the Lorenz equation. Comparing these solutions to the physical solution visualized in Figures 3-1–3-3, we see that the periodic solutions are significantly more limited in their extent of exploration.

Because periodic solutions have a more limited extent in the phase space, their statistics are different from physical solutions. Here we illustrate the difference using the mean of two quantities of interest

$$J_1(u) = u_3 \quad \text{and} \quad J_2(u) = e^{-\frac{u_3^2}{2}}. \quad (3.5)$$

Table 3.1 shows these quantities of interest averaged over the three periodic solutions shown in Figure 3-4, compared against those averaged over a physical solution. Here, Periodic #1, #2, and #3 correspond to the solid, dashed, and dotted lines, respec-

tively. We observe from the table that the mean of J_1 over the periodic solutions is different but comparable to the mean over physical solutions. The mean of J_2 over the periodic solutions, on the other hand, is orders of magnitude different from that of the physical solutions. These differences disqualify the periodic solutions from being physical.

	Periodic #1	Periodic #2	Periodic #3	Physical solutions
J_1	23.05	23.19	23.37	23.67
J_2	4×10^{-35}	7×10^{-28}	3×10^{-22}	1.58×10^{-05}

Table 3.1: Comparison of statistics computed from periodic solutions of the Lorenz equation with the statistics computed from physical solutions

While periodic solutions are generally difficult to find, the Lorenz equation has a special feature that makes the task significantly easier. In a typical solution to the Lorenz equation, the $u_3(t)$ component oscillates in a pattern that appears neither regular nor random. Lorenz observed that the height of one peak in the oscillation can predict the height of the next peak. He quantified his observation by the Lorenz map, as shown in Figure 3-5.

The Lorenz map provides us a tool to find as many periodic solutions as we want. By intersecting the map with a diagonal line, we can find a local maximum of $u_3(t)$ for which the next maximum is almost the same value. We then look up the values of $u_1(t)$ and $u_2(t)$ when $u_3(t)$ achieves this maximum. This gives us an initial condition starting from which the solution is nearly periodic. We can similarly intersect the second iterate of the Lorenz map (right plot of Figure 3-5), and the third iterate, etc, with a diagonal line, to find increasingly complex periodic solutions.

This Lorenz map is more than a tool to study the Lorenz equation. It is a chaotic dynamical system all by itself. Unlike the Lorenz equation, which is a continuous-time dynamical system in three dimensions, the Lorenz map is a discrete-time dynamical system in one dimension. It exhibits the same sensitivity to initial condition that characterizes the Lorenz equation. One can readily observe in the right plot of Figure 3-5 that a small perturbation in the x-axis can lead to a large change in the y-axis.




Figure 3-5: L: the Lorenz map. The x-axis is the n -th local maximum of $u_3(t)$ over a long solution; the y-axis is the $(n + 1)$ -th local maximum of $u_3(t)$. The intersection of this curve with the dashed line ($y = x$) indicates the initial condition for the solid line in Figure 3-4. R: the Lorenz map iterated twice. The x-axis is the n -th local maximum of $u_3(t)$; the y-axis is the $(n + 2)$ -th maximum. The intersections with the diagonal dashed line indicate the initial conditions for both the dotted and dashed lines in Figure 3-4.

This sensitivity grows exponentially for iterates of this map. A solution to the Lorenz map can be obtained by extracting the consecutive local maxima of a solution to the Lorenz equation. If we extract from a physical solution to the Lorenz equation, we obtain a physical solution of the Lorenz map. It will visit the interval between 29 and 49 with varying, well-defined frequencies. By contrast, if one extracts a solution to the Lorenz map from a periodic solution to the Lorenz equation, the solution will visit only a discrete set of points. It is thus a periodic, nonphysical solution. What we learned about the Lorenz equation could have all been learned from the Lorenz map.

Having discussed periodic solutions in this section, we move to a second type of nonphysical solutions. This type is more difficult to find and study than the periodic solutions. To make it easier, we switch our example to a one-dimensional, discrete-time dynamical systems like the Lorenz map. Because the Lorenz map lacks a closed form, we construct, in the next subsection, a one-dimensional discrete-time dynamical




Figure 3-6: The tent map φ .

system with a closed form, one that qualitatively resembles the Lorenz map. This map will help us, in section 3.3.4, to study a more insidious type of nonphysical solutions that hides in shadowing solutions.

3.3.3 Tent map: periodic and physical solutions

The tent map is qualitatively similar to the Lorenz map,

$$\varphi(x) := \begin{cases} 2x & x < 1 \\ 4 - 2x & 2 \geq x \geq 1 \end{cases}. \quad (3.6)$$

Due to its simple analytical form, the properties of the tent map have been studied extensively [86][279]. Figure 3-6 shows the tent map.

The tent map φ is chaotic because a trajectory x_0, x_1, \dots satisfying $x_{i+1} = \varphi(x_i)$ exhibits exponential sensitivity to initial condition. An infinitesimal perturbation of absolute value δx , applied to an initial condition x_0 , generates a trajectory that is $2\delta x$ away from x_1 , $4\delta x$ from x_2 , $8\delta x$ from x_3 and so on. This exponential divergence of two trajectories that start infinitesimally apart is the butterfly effect that characterizes chaotic dynamics.

It is easy to find periodic solutions for the tent map. $\frac{4}{3}$ maps to itself; $\frac{4}{5}$ and $\frac{8}{5}$ map to each other; $\frac{4}{7}$, $\frac{8}{7}$, and $\frac{12}{7}$ map circularly. In fact, any rational number evolves into periodic solutions that visit only a finite set of rational numbers with the same denominator. This is compatible with ergodic theory because all rational numbers in $[0, 2]$ comprise a subset of Lebesgue measure zero. That is, we would have zero likelihood of getting a rational number if we sampled the Lebesgue measure (the uniform distribution) on $[0, 2]$. Instead, with a 100% probability, one would get an irrational number that, when iterated under the tent map, leads to a physical trajectory that distributes uniformly on the interval $[0, 2]$.

To understand why a physical solution visits the interval $[0, 2]$ at a uniform frequency, it is helpful to view the tent map from a different perspective. For a randomly chosen $x_0 \in [0, 2]$, we can represent it in binary form:

$$x_0 = \sum_{j=0}^{\infty} \frac{x_{0,j}}{2^j} \quad (3.7)$$

where $x_{0,0} \in \{0, 1\}$ is the integer component of x_0 and $x_{0,j} \in \{0, 1\}$ is the j th bit after the binary point. Let $x_{i+1} = \varphi(x_i)$, then it is straightforward to verify from the definition of the map that its binary representation

$$x_i = \sum_{j=0}^{\infty} \frac{x_{i,j}}{2^j} \quad (3.8)$$

satisfies

$$x_{i+1,j} = x_{i,0} \underline{\vee} x_{i,j+1}, \quad (3.9)$$

where $\underline{\vee}$ is the xor operator. To see why, note that multiplication by 2 is a left-shift operator in binary, and subtraction from 4 flips every bit after the binary point. If x_0 is chosen uniformly in $[0, 2]$, then each bit $x_{0,j}$, $j = 0, 1, \dots$ of x_0 has equal probability of being 0 or 1, and each bit is independent of other bits. It follows from Eq. 3.9 that each bit $x_{i,j}$, $j = 0, 1, \dots$ of x_i , $i = 1, 2, \dots$ has equal probability of being 0 or 1, and each bit is still independent of the other bits. As the map iterates starting from almost any x_0 , a physical solution explores the entire interval $[0, 2]$ uniformly.

3.3.4 Nonphysical solutions type II: Quasi-physical solutions

The simplicity of the tent map, as well as its binary form (Eq. 3.9), enables us to study nonphysical solutions that are not periodic. As in the last section, consider an

$$x_0 = \sum_{j=0}^{\infty} \frac{x_{0,j}}{2^j} \quad (3.10)$$

in which the bits $x_{0,j}$ are not independent of each other. Instead, suppose each bit is more likely to be identical to the previous bit than to be different. That is, $x_{0,j+1} = x_{0,j}$ with probability $p > \frac{1}{2}$ for $j = 0, 1, 2, \dots$. Then, by Eq. 3.9, the bits of every $x_i, i = 1, 2, \dots$ follow the same pattern, namely, each bit repeats the previous bit with probability p . Moreover, consider Eq. 3.9 for $j = 0$ and any i :

$$x_{i+1,0} = x_{i,0} \vee x_{i,1}. \quad (3.11)$$

Because $x_{i,0} = x_{i,1}$ with probability $p > \frac{1}{2}$, $x_{i+1,0} = 0$ with probability $p > \frac{1}{2}$. Starting from x_1 , this solution visits $[0, 1]$ with probability $p > \frac{1}{2}$. Instead of visiting $[0, 1]$ and $[1, 2]$ with equal probability, as a physical solution does, this solution favors $[0, 1]$. Since it provably visits the interval $[0, 2]$ with a different frequency from that of a physical solution, this is a nonphysical solution.

If $p = 1$, the bits of x_0 are either all zeros or all ones, which is the same as all zeros in mod 2 arithmetic. So, all further iterates when x_0 is 0, are 0, and this is a trivial nonphysical solution of type 1. Now when $p > \frac{1}{2}$ but strictly less than 1, the solutions we just constructed are both nonphysical and aperiodic. The bits of x_0 , though correlated with each other, still can exhibit an infinite variety of patterns. This implies that the solution may not eventually converge to any fixed point, nor follow any strict period. As the map iterates, the nonrepetitive bit patterns shift towards more significant digits, and the solution visits an infinite set of points. Nevertheless, the solutions observed along a trajectory also do not conform to a uniform distribution on $[0, 2]$ since they preferentially visit the first half of this interval. We call such aperiodic nonphysical solutions “quasi-physical” solutions.




Figure 3-7: Empirical probability distribution of long solutions (a billion steps) starting from four points whose binary digits have probability p of repeating the previous digit. The solution shown in the top-left plot starts from a point with $p = 0.5001$; top-right: $p = 0.51$; bottom-left: $p = 0.55$; bottom-right: $p = 0.9$.

Figure 3-7 shows the empirical density functions of four such quasi-physical solutions. When $p = 0.5001$, the statistical distribution of the quasi-physical solution is approximately uniform. Recall that a physical solution explores $[0, 2]$ uniformly. In this case, the quasi-physical solution has very similar statistical behavior as a physical solution. When $p = 0.51$, the empirical distribution becomes “hairy”. An apparently fractal pattern emerges. This fractal pattern further amplifies when $p = 0.55$. Meanwhile, the density on the left sub-interval, $[0, 1]$, becomes obviously higher than the density on the right sub-interval, $[1, 2]$. This is consistent with our theoretical analysis at the beginning of this subsection. When $p = 0.9$, the fractal pattern is so intensified that most of the solution seems to concentrate in a collection of tiny intervals. These plots expose the diversity of quasi-physical solutions. Their statistical distribution can be either very similar or completely different from that of physical solutions.

So far we have constructed one class of quasi-physical solutions. It is noteworthy that there are infinite ways to construct quasi-physical solutions. In the binary representation of the initial condition x_0 , any statistical deviation from equal probability of 0 or 1, or any statistical dependence among digits would lead to quasi-physical solutions. One could, for example, construct an x_0 in which a bit is more likely to be 1 than 0 only if it follows two consecutive 0’s. Such an x_0 would lead to a nonphysical solution. Its statistical distribution would differ from any of the plots in Figure 3-7. Nevertheless, as can be observed in Figure 3-8, it shows a remarkable resemblance in its “hairiness”; some kind of fractal pattern emerges from the distribution. A fractal distribution appears to be a signature of quasi-physical solutions.

We have only seen quasi-physical solutions for the tent map. It is difficult to analytically construct quasi-physical solutions to the Lorenz equation and other, more complex, governing equations which typically produce chaotic solutions. Nevertheless, the similarity between the tent map and the Lorenz map [279], shown in Figure 3-5, suggests that quasi-physical solutions may exist for the Lorenz map, and by extension, the Lorenz equation. It is then natural to conjecture that such aperiodic, nonphysical, quasi-physical solutions exist in general for chaotic dynamical systems.

These quasi-physical solutions raise doubts over the usefulness of shadowing in




Figure 3-8: Empirical distribution function of other quasi-physical solutions. The solution shown in the left plot starts from an initial condition whose bits are independent and have probability 0.6 of being 0. The solution shown in the right plot starts from an initial condition whose bits have a probability of 0.6 of being 1 only following two consecutive 1's; otherwise a bit is 0 or 1 with equal probability.

some applications. For example, even if a numerical solution is shadowed by a solution to the true governing physics, is the shadowing solution physical or quasi-physical? At first glance, this may seem to be a nonissue because almost all solutions are physical. It seems reasonable to argue that because the set of all nonphysical solutions is Lebesgue measure-zero, the probability of finding a nonphysical solution through shadowing is zero percent. Such an argument, as we show in the next section, is wrong. The probability of finding a nonphysical solution through shadowing can be, instead of zero percent, one hundred percent.




Figure 3-9: L: The scaled tent map $\hat{\varphi}_s$ at different values of s between 0 and 1. R: sensitivity to small perturbation in the governing equation. The y-axis shows the absolute value of the difference between two solutions, one satisfying Eq. 3.12, one for $s = 0$ and the other for $s = 10^{-5}$. The x-axis shows the iteration number. The initial condition is at $x_0 = \pi/2$.

3.4 Are shadowing solutions physical?

3.4.1 Example of a shadowing solution for the tent map

To illustrate the concept of shadowing, consider the tent map, defined in Eq. 3.6, and a scaled version of the map, defined by

$$\hat{\varphi}_s(x) := \begin{cases} 2x & x < 1 + s \\ 4(1 + s) - 2x & x \geq 1 + s \end{cases} \quad (3.12)$$

where $s \ll 1$. Note that a small change in s can lead to drastic differences between solutions starting from the same initial condition. As demonstrated in Figure 3-9, this sensitivity to small perturbations reflects the chaotic nature of the governing equation.

We can avoid this sensitivity to the governing equation using a coordinated perturbation to the initial condition. Consider two solutions satisfying Eq. 3.12 at different values of s . Instead of starting from the same initial condition, these two solutions

start from $x_0(1+s)$ with the same x_0 but their respective values of s . It can be shown that the solution of these two equations would be $x_i(1+s), i = 0, 1, \dots$ with their respective values of s , where $x_{i+1} = \varphi(x_i)$ satisfies the original tent map (Eq. 3.6). When the values of s are similar between the two solutions, this pair of solutions stays uniformly close to each other forever. Such a true solution of the governing equation that stays close to a given perturbed solution over a long time, is known as a *shadowing* solution. In this example, the perturbed solution satisfies a slightly different governing equation.

For every solution $\hat{x}_i, i = 0, \dots$ satisfying Eq. 3.12, there is a shadowing solution satisfying Eq. 3.6: $x_i := \hat{x}_i/(1+s)$. The map,

$$\hat{h}_s(x) := x/(1+s) \tag{3.13}$$

is known as a *conjugacy* between the two maps φ and $\hat{\varphi}_s$ because it satisfies

$$\varphi(\hat{h}_s(x)) = \hat{h}_s(\hat{\varphi}_s(x)), \quad \forall x \tag{3.14}$$

or equivalently, $\varphi \circ \hat{h}_s \equiv \hat{h}_s \circ \hat{\varphi}_s$. Such conjugacy maps can generally help us construct shadowing solutions. For every solution satisfying $\hat{\varphi}_s$, h_s maps it to a shadowing solution satisfying φ because iterating Eq. 3.14, we get $\varphi^n \circ \hat{h}_s(x) = \hat{h}_s \circ \hat{\varphi}_s^n(x)$, for $n = 1, 2, 3 \dots$, where f^n stands for the function composition of f n -times.

Is the shadowing solution a physical solution? In this example, the answer is almost surely positive. We can demonstrate that almost any solution of the scaled tent map (Eq. 3.12) is uniformly distributed in $[0, 2(1+s)]$ – we can repeat our analysis in section 3.3.3 but represent our initial condition as $x_0 = (1+s) \sum_{j=0}^{\infty} x_{0,j}/2^j$. Its shadowing solution, which satisfies the original tent map (Eq. 3.6), can be obtained through the conjugacy (Eq. 3.13). Thus, the distribution of the shadowing solution can be obtained by mapping a uniform distribution in $[0, 2(1+s)]$ through the conjugacy map. This leads to a uniform distribution in $[0, 2]$, which is precisely the distribution of a physical solution of the tent map.

This simple example is useful in illustrating the concept of shadowing and the

utility of the conjugacy map. But it is rather special in that the shadowing solution is almost always physical. Our next example introduces a tilted version of the tent map. Although the tilting perturbation to the tent map seems as simple as the scaling perturbation, the shadowing solution, as we will see, is almost always a quasi-physical solution.

3.4.2 An example of quasi-physical shadowing solution

The tilted tent map is defined as

$$\tilde{\varphi}_s(x) := \begin{cases} \frac{2}{1+s}x & x < 1+s \\ \frac{2}{1-s}(2-x) & x \geq 1+s \end{cases} \quad (3.15)$$

When $s = 0$, this map is identical to the tent map (Eq. 3.6). Let $\tilde{x}_{i+1} = \tilde{\varphi}_s(\tilde{x}_i)$, $i = 0, 1, \dots$ be a solution of $\tilde{\varphi}_s$ that is also a perturbed solution of φ . As we explained in the previous subsection, the corresponding shadowing solution, which solves φ , can be found through a conjugacy map \tilde{h}_s that connects the solutions of φ and $\tilde{\varphi}_s$. If $\varphi \circ \tilde{h}_s \equiv \tilde{h}_s \circ \tilde{\varphi}_s$, then $x_i = \tilde{h}_s(\tilde{x}_i)$, $i = 0, 1, \dots$ is the shadowing solution that satisfies $x_{i+1} = \varphi(x_i)$, corresponding to the perturbed solution \tilde{x}_i , $i = 0, 1, \dots$.

The conjugacy map, although more complex than the one in the last subsection, has the following closed form:

$$\tilde{h}_s(x) = \sum_{j=0}^{\infty} \frac{\xi_{s,j}(x)}{2^j}, \quad (3.16)$$

where

$$\xi_{s,j}(x) := \begin{cases} 0 & j = 0, x < 1+s \\ 1 & j = 0, x \geq 1+s \\ \xi_{s,j-1}(x) & j > 0, \tilde{\varphi}_s^j(x) < 1+s \\ 1 - \xi_{s,j-1}(x) & j > 0, \tilde{\varphi}_s^j(x) \geq 1+s. \end{cases} \quad (3.17)$$

In the above expression, $\tilde{\varphi}_s^j$ refers to the j -time composition of $\tilde{\varphi}_s$. That is, if $\tilde{x}_{j+1} =$




Figure 3-10: L: the tilted tent map Eq. 3.15 for $s = 0, 0.1, 0.2, 0.3, 0.4$, and 0.5 . R: the conjugacy \tilde{h}_s between the tent map (Eq. 3.6) and the tilted tent map (Eq. 3.15), evaluated using Eq. 3.16-3.17, for the same set of s as the left plot.

$\tilde{\varphi}_s(\tilde{x}_j), j = 0, 1, \dots$, is a solution, $\tilde{\varphi}_s^i(\tilde{x}_j) = \tilde{x}_{j+i}$, $i, j = 0, 1, \dots$.

To see why \tilde{h}_s constructed in Eq. 3.16 satisfies the definition of a conjugacy map – $\tilde{h}_s(\tilde{\varphi}_s(x)) = \varphi(\tilde{h}_s(x))$ for all x – we need to analyze two cases: $x < 1 + s$ and $x \geq 1 + s$. When $x < 1 + s$, $\xi_{s,0}(x) = 0$; thus $\tilde{h}_s(x) = \sum_{j=1}^{\infty} \frac{\xi_{s,j}(x)}{2^j}$. Inside this infinite series, $\xi_{s,1}(x) = \xi_{s,0}(x) = 0$ if $\tilde{\varphi}_s(x) < 1 + s$, or $\xi_{s,1}(x) = 1 - \xi_{s,0}(x) = 1$ if $\tilde{\varphi}_s(x) \geq 1 + s$. Thus, $\xi_{s,1}(x) = \xi_{s,0}(\tilde{\varphi}_s(x))$ according to the definition of $\xi_{s,0}$. Using this as the base case, one can inductively verify that $\xi_{s,j+1}(x) = \xi_{s,j}(\tilde{\varphi}_s(x))$ for all $j > 0$, using just the definitions of $\xi_{s,j}$ and $\xi_{s,j+1}$. Therefore, $\tilde{h}_s(x) = \frac{1}{2} \sum_{j=0}^{\infty} \frac{\xi_{s,j}(\tilde{\varphi}_s(x))}{2^j} = \frac{1}{2} \tilde{h}_s(\tilde{\varphi}_s(x))$. On the other hand, because $\tilde{h}_s(x) = \sum_{j=1}^{\infty} \frac{\xi_{s,j}(x)}{2^j} \leq 1$, $\varphi(\tilde{h}_s(x)) = 2\tilde{h}_s(x)$ according to the definition of φ . Thus, $\varphi(\tilde{h}_s(x)) = \tilde{h}_s(\tilde{\varphi}_s(x))$ holds when $x < 1 + s$.

When $x \geq 1 + s$, $\xi_{s,0}(x) = 1$; thus $\tilde{h}_s(x) = 1 + \sum_{j=1}^{\infty} \frac{\xi_{s,j}(x)}{2^j} = 2 - \sum_{j=1}^{\infty} \frac{1 - \xi_{s,j}(x)}{2^j}$. Inside this infinite series, $1 - \xi_{s,1}(x) = 1 - \xi_{s,0}(x) = 0$ if $\tilde{\varphi}_s(x) < 1 + s$, or $1 - \xi_{s,1}(x) = \xi_{s,0}(x) = 1$ if $\tilde{\varphi}_s(x) \geq 1 + s$. Thus, $1 - \xi_{s,1}(x) = \xi_{s,0}(\tilde{\varphi}_s(x))$ according to the definition of $\xi_{s,0}$. Using this as the base case, one can inductively verify that $1 - \xi_{s,j+1}(x) = \xi_{s,j}(\tilde{\varphi}_s(x))$ for all $j > 0$, using just the definitions of $\xi_{s,j}$ and $\xi_{s,j+1}$. Therefore, $\tilde{h}_s(x) = 2 - \frac{1}{2} \sum_{j=0}^{\infty} \frac{\xi_{s,j}(\tilde{\varphi}_s(x))}{2^j} = 2 - \frac{1}{2} \tilde{h}_s(\tilde{\varphi}_s(x))$. On the other hand, because $\tilde{h}_s(x) = 1 + \sum_{j=1}^{\infty} \frac{\xi_{s,j}(x)}{2^j} \geq 1$, $\varphi(\tilde{h}_s(x)) = 4 - 2\tilde{h}_s(x)$ according to the definition of φ .




Figure 3-11: L: the density of a trajectory satisfying Eq. 3.6 that shadows a random trajectory satisfying Eq. 3.15 for $s = 0.1$. R: the density of a trajectory satisfying Eq. 3.6 that shadows a random trajectory satisfying Eq. 3.15 for $s = 0.5$.

Thus, $\varphi(\tilde{h}_s(x)) = \tilde{h}_s(\tilde{\varphi}_s(x))$ also holds when $x \geq 1 + s$.

Figure 3-10 shows how fractal the conjugacy map is. This fractal conjugacy map transforms a uniform distribution in $[0, 2]$ into a fractal distribution, shown in Figure 3-11, similar to the ones plotted in section 3.3.4. In fact, we will show that the fractal distribution obtained by mapping a uniform distribution through \tilde{h}_s is exactly the family of distributions shown in Figure 3-7. We will also show that a physical solution of the tilted tent map (Eq. 3.15) is uniformly distributed in $[0, 2]$. Thus, for almost any physical solution of the tilted tent map, a shadowing solution of the original tent map, obtained through the conjugate map \tilde{h}_s , has a fractal distribution. Such a shadowing solution is therefore a quasi-physical solution of the tent map.

We first show that a physical solution of the tilted tent map (3.15) is uniformly distributed in $[0, 2]$, for any $0 \leq s < 1$. For this, we only need to show that the uniform distribution is stationary under the tilted tent map. The function $\tilde{\varphi}_s$ stretches an infinitesimal region around each x by the absolute value of the derivative at x ; hence, a uniform density on an infinitesimal region around $\tilde{\varphi}_s(x)$ is reduced by the same factor. In other words, let $\tilde{\rho}(x)$ be the map of a uniform density on $[0, 2]$, through $\tilde{\varphi}_s$.

Then, $\tilde{\rho}(\tilde{\varphi}_s(x)) = 0.5/|\tilde{\varphi}'_s(x)| + 0.5/|\tilde{\varphi}'_s(x')|$, where x and x' lie on two sides of $1 + s$ and $\tilde{\varphi}_s(x) = \tilde{\varphi}_s(x')$; the two terms express the conservation of probability from the two pre-images of $\tilde{\varphi}_s$. Substituting the piecewise constant derivative of $\tilde{\varphi}_s$ on the two intervals, $[0, 1 + s)$ and $[1 + s, 2]$, we have $\tilde{\rho}(\tilde{\varphi}_s(x)) = 0.25(1 + s) + 0.25(1 - s) = 0.5$, at all x . Hence, a uniform density of 0.5 is unaltered by mapping through $\tilde{\varphi}_s(x)$.

We now show that the conjugacy \tilde{h}_s maps a uniform distribution to a fractal. Here we use the closed forms Eqs. 3.16-3.17. If x is a random number drawn uniformly in $[0, 2]$, it has $\frac{1+s}{2}$ probability of being less than $1 + s$. Thus, $\xi_{s,0}(x) = 0$ with probability $\frac{1+s}{2}$. This means, according to Eq. 3.16, $\tilde{h}_s(x) < 1$ with probability $\frac{1+s}{2}$: $\tilde{h}_s(x)$ is more likely to lie in the left half of the domain $[0, 2]$. Furthermore, for $j \geq 0$, $\xi_{s,j+1}(x) = \xi_{s,j}(x)$ with probability $\frac{1+s}{2}$ since each $\tilde{\varphi}_s^j(x)$ is sampled from the uniform distribution on $[0, 2]$. This probability has direct implication on $\tilde{h}_s(x)$ since, again by Eq. (3.16), $\xi_{s,j}$ is the j th bit in the binary representation of \tilde{h}_s – each bit repeats the previous one with probability $\frac{1+s}{2}$. For a uniformly random x , $\tilde{h}_s(x)$ is exactly the kind of initial condition we used to construct the quasi-physical solution in section 3.3.4, with $p = \frac{1+s}{2}$. Since this analysis holds for any x sampled uniformly on $[0, 2]$, we can conclude that the shadowing solution is almost surely nonphysical. Let $x_{s,i}, i = 0, 1, \dots$ be a solution to the tilted tent map $\tilde{\varphi}_s$ with $x_{s,0}$ chosen randomly in $[0, 1 + s]$. Then, with a hundred percent probability, it is a physical solution that uniformly visits the domain $[0, 2]$. Its shadowing solution $\tilde{h}_s(x_{s,i}), i = 0, 1, \dots$, however, is a quasi-physical solution, also with a hundred percent probability. It explores $[0, 2]$ at a nonuniform frequency with a fractal probability distribution.

3.4.3 Are general shadowing solutions physical?

In the previous subsection, we examined a tilting perturbation to the tent map. We derived the fractal conjugacy map, and showed that a typical shadowing trajectory is exactly the quasi-physical solution we analyzed in section 3.3.4. This example contrasts with section 3.4.1, in which the shadowing solutions of a different perturbation, the scaling perturbation, are physical solutions. Which situation is more typical? When we make other types of perturbations, do we expect to observe physical or




Figure 3-12: The “squashed” tent map (Eq. 3.18) for $s = 0.01, 0.05, 0.2$, and 0.5 .

nonphysical shadowing solutions?

This section attempts to answer this question by investigating several other perturbations, the first of which scales the height, but not the width, of the tent map. The resulting “squashed” tent map, illustrated in Figure 3-12, has the closed form

$$\hat{\varphi}_s^{\text{sq}}(x) := \begin{cases} (2-s)x & x < 1 \\ (2-s)(2-x) & 2 > x \geq 1. \end{cases} \quad (3.18)$$

In Figure 3-13, we show side-by-side the probability distributions generated by observing a long solution (of length 10 billion) of $\hat{\varphi}_s^{\text{sq}}$ (left) and its corresponding shadowing solution (right) at two different values of s . The left column showing the probability distribution of the shadowing solution is reminiscent of the hairy probability distributions of the quasi-solutions of the tent map (section 3.3.4). The reader is referred to the Supplementary Material section 3.6.2 for the computational method used in this chapter, for the shadowing solutions. Note that, unlike the tilted tent maps, the family of squashed tent maps do not have the uniform distribution as their SRB distribution, but have a different regular probability distribution as indicated on the right column of Figure 3-13. On the other hand, the shadowing solution is distributed neither like a physical solution of the original tent map nor of the squashed




Figure 3-13: The left column shows the empirical probability distribution of the shadowing solution for the squashed tent map at $s = 0.01$ (top) and at $s = 0.2$ (bottom). The right column shows the empirical probability distribution of a physical solution at the same two values of $s = 0.01$ (top) and $s = 0.2$ (bottom).




Figure 3-14: The pinched tent map (Eq. 3.19) for $s = 0.01$ (blue), $s = 0.05$ (orange), $s = 0.2$, (green) and $s = 0.5$ (red).

tent map, as suggested by its fractal-like probability distribution on the left column of Figure 3-13.

Next we consider a second perturbation of the tent map, the *pinched* tent map, which has the following closed form

$$\hat{\varphi}_s^p(x) = \begin{cases} \frac{4x}{1 + s + \sqrt{(1 + s)^2 - 4sx}}, & x < 1 \\ \frac{4(2 - x)}{1 + s + \sqrt{(1 + s)^2 - 4s(2 - x)}}, & 2 \leq x \leq 1. \end{cases} \quad (3.19)$$

Again, at $s = 0$, the original tent map is recovered, and at other values of s , the tent map is “pinched” by perturbations that are symmetric around $x = 1$, and zero at the end points. Figure 3-14 shows the pinched tent map at different s values. In Figure 3-15, we show the corresponding plots of the physical and shadowing distributions for the pinched tent map. From the right column, which shows the probability density of a long physical solution, we can see that the pinching perturbation changes the uniform density of the original tent map to a linearly varying density. The more pronounced the perturbation – the higher the value of s – the steeper the slope. The shadowing distribution, in this case as well, appears fractal. Once again, with probability one – or, for any randomly chosen initial condition of the solution of the

pinched tent map – a physical solution is shadowed by a quasi-physical solution. We consider yet another perturbation of the tent map with the closed form

$$\hat{\varphi}_s^w(x) = \begin{cases} \frac{4x}{1 + s + \sqrt{(1 + s)^2 - 4sx}}, & x < 1 \\ \frac{4(2 - x)}{1 - s + \sqrt{(1 - s)^2 - 4s(2 - x)}}, & 2 \leq x \leq 1. \end{cases} \quad (3.20)$$

The above map, called the wave tent map, is illustrated in Figure 3-16. To construct the pinched tent map (Eq. 3.19), we perturbed the original tent map in such a way that the perturbations on the two halves of the interval $[0, 2]$ are mirror images of each other about the $x = 1$ line. In the wave tent map, the perturbations on the two halves are a reflection of each other about the x -axis. The wave perturbation does not noticeably alter the SRB density of the original tent map. As we show in Figure 3-17 (right), a solution of the wave tent map, starting from almost any point on $[0, 2]$ – a physical solution – has a density that looks almost identical to a uniform density of 0.5. On the left column of Figure 3-17, we observe that the distributions of the shadowing solutions corresponding to the physical solutions on the right once again appear fractal. Hence, we can once again conclude that almost every solution, which is physical, has a corresponding shadowing solution that is nonphysical. In this section, we have shown that the same conclusion holds for several different perturbations of the tent map. The physical solutions were distributed nonuniformly and differently from one another. However, we observed a commonality in the different perturbations: the shadowing solutions corresponding to almost every physical solution had a fractal-like probability distribution that did not resemble the distributions of the physical solutions in any case. Thus, we have shown significant evidence, through analytical constructions of perturbed tent maps in this section and section 3.3.4 that shadowing solutions can almost surely be nonphysical.




Figure 3-15: The left column shows the empirical distribution of the shadowing solution for the pinched tent map at $s = 0.05$ (top) and at $s = 0.2$ (bottom). The right column shows the probability distribution of a physical solution at the same two values of $s = 0.05$ (top) and $s = 0.2$ (bottom).




Figure 3-16: The wave tent map (Eq. 3.20) for $s = 0.01, 0.05, 0.2$, and 0.5 .

3.4.4 What about the Lorenz equation?

In order to analyze the behavior of the shadowing solutions of the Lorenz'63 system of equations, we first make a closed form approximation of the Lorenz map (see section 3.6.1 of the supplementary material). Except at rare parameter values, the Lorenz system does not have solutions that shadow perturbed orbits for all time, but rather only for a finite time [157]. Hence, the Lorenz map also does not have infinitely long shadowing solutions. This implies that a numerical solution of the Lorenz map approximates a true solution, as accurately as desired, only for a finite time.

A closed form approximation is constructed by regression performed using long solutions of the Lorenz map (Figure 3-5). The Lorenz map approximation obtained this way is illustrated in Figure 3-18, for variations of the parameters around their standard values of $\rho = 28$, $\sigma = 10$, and $\beta = 8/3$. Using the approximate Lorenz map, we calculate next the shadowing solutions. The empirical distribution computed from long shadowing solutions, of length 10 billion, are shown in the left column of Figure 3-19 at three different sets of parameter values, which are also different from the standard values of $\rho = 28$, $\sigma = 10$, and $\beta = 8/3$. In each row, one of the parameters is perturbed from the standard values. On the right column, we plot the distribution of physical solutions, the SRB distribution, at the set of parameters corresponding to




Figure 3-17: The left column shows the empirical distribution of the shadowing solution for the wave tent map (Eq. 3.20) at $s = 0.05$ (top) and at $s = 0.2$ (bottom). The right column shows the probability distribution of a physical solution at the same two values of $s = 0.05$ (top) and $s = 0.2$ (bottom).




Figure 3-18: The left, center, and right plots show the effect of the parameters σ , ρ , and β on the Lorenz map, respectively. In the left plot, the blue and orange lines represent $\sigma = 10$ and 12 , respectively, while $\rho = 28$ and $\beta = 8/3$. In the center plot, the blue and orange lines represent $\rho = 28$ and 30 , respectively, while $\sigma = 10$ and $\beta = 8/3$. In the right plot, the blue and orange lines represent $\sigma = 8/3$ and $10/3$ respectively, while $\sigma = 10$ and $\rho = 28$.

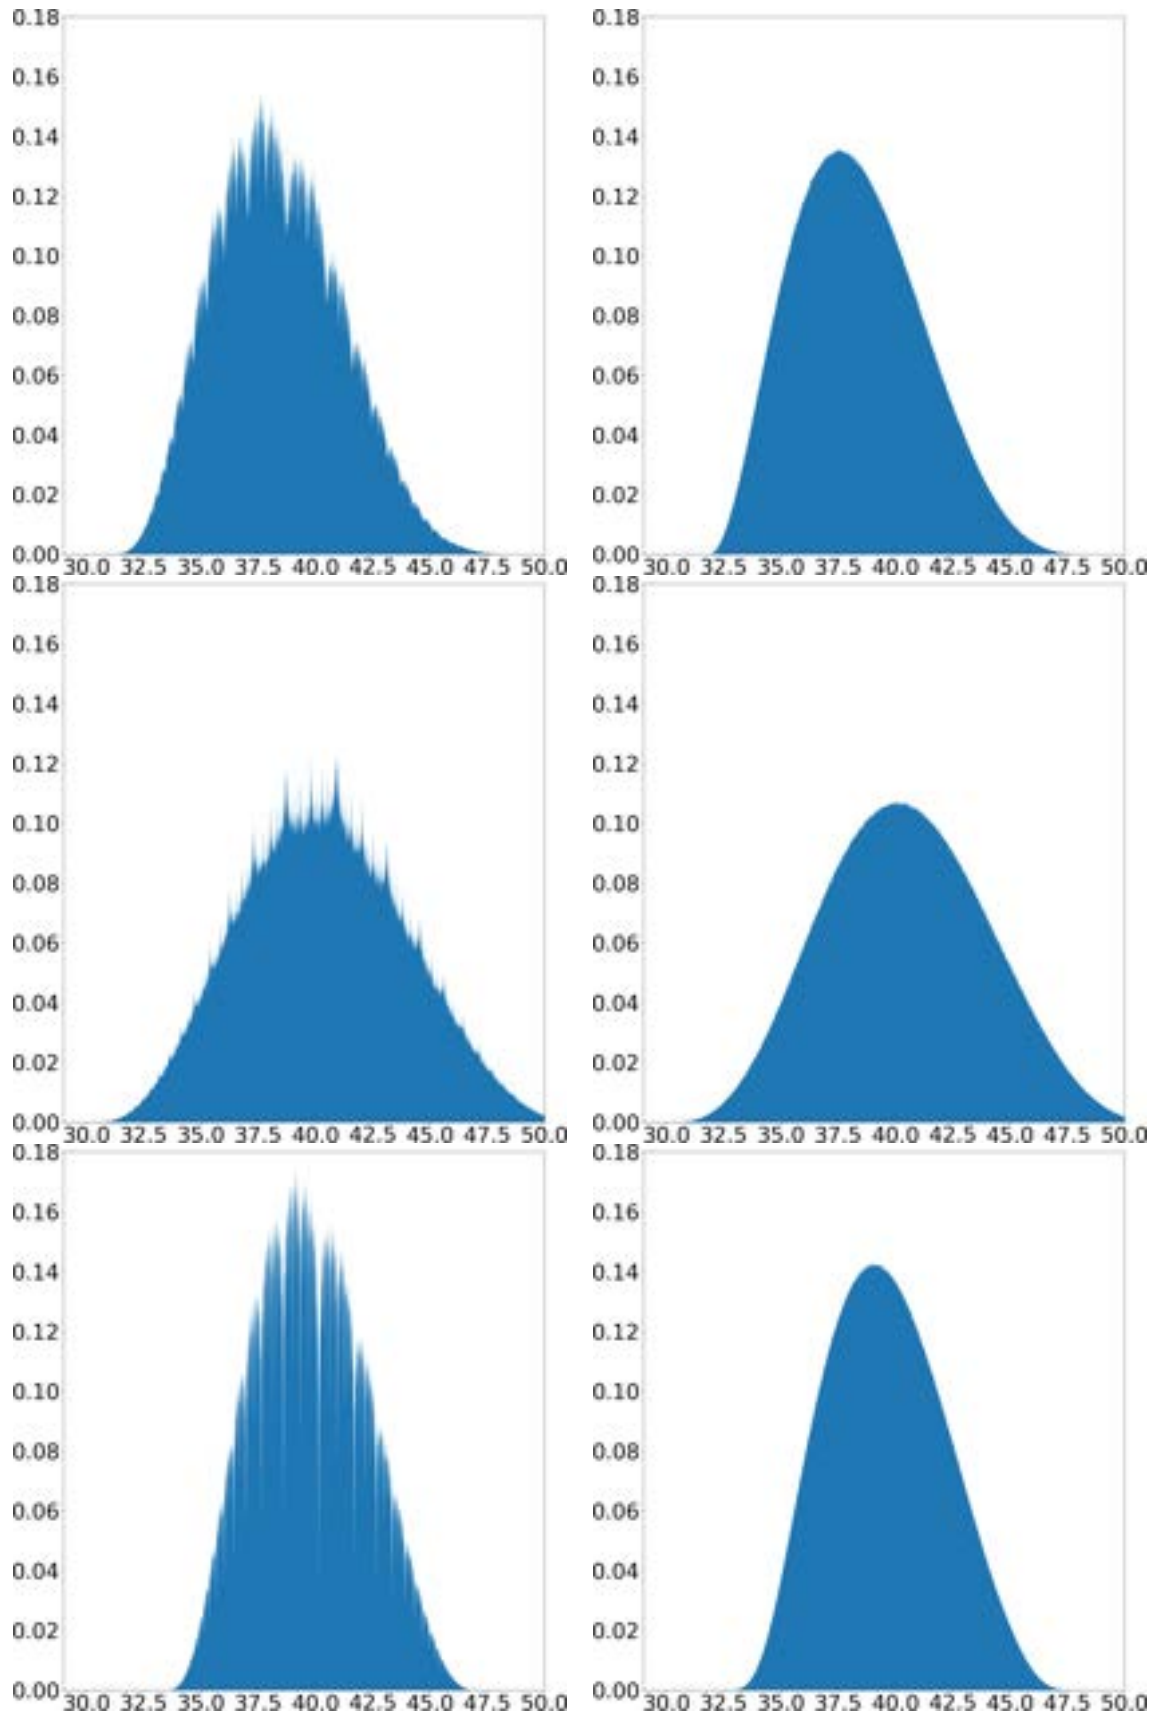


Figure 3-19: The left column shows the empirical distribution of the shadowing solution, and the right column shows the physical distribution of the Lorenz map at the following sets of parameters: top row: $\sigma = 15, \rho = 28, \beta = 8/3, \rho = 29$, middle row: $\sigma = 10, \rho = 30, \beta = 10/3$, and bottom row: $\sigma = 10, \rho = 28, \beta = 10/3$.

each row. We again notice that, while the physical distributions appear smooth, the shadowing solutions clearly exhibit roughness. Thus, we draw the same conclusion in the Lorenz map as in the various perturbations of the tent map treated in section 3.4.3: shadowing solutions do not have the same distribution as a typically observed solution of the governing equation. In other words, shadowing solutions are nonphysical.

3.5 Discussion and outlook

Through rigorous counterexamples, we show that shadowing can lead to nonphysical solutions. This has a troubling implication for numerical simulations of chaotic governing equations. Through the examples in this chapter, it is clear that a numerical solution, satisfying a governing equation perturbed due to numerics, is not expected to be shadowed by a physical solution of the real governing physics. What is, if any, the relation of a numerical solution, such as DNS of turbulent flows, to the true physics? For an $\mathcal{O}(10^{-18})$ perturbation (due to numerical error), we may get completely different statistics from a true, physical solution. How then can we trust numerical solutions when they are not guaranteed to share the long-term statistical behavior of the governing equation? We conclude with an example that illustrates a counterintuitive feature of some chaotic systems: a small perturbation to the governing equation can significantly change the statistics/long-term behavior of its physical solutions. We construct yet another perturbed tent map – the *plucked* tent map – in which an oscillatory perturbation is introduced (see Supplementary Material section 3.6.3 for the map equation). In Figure 3-20, we show that the magnitude of the oscillatory perturbation is controlled by parameters s and n , and its frequency is controlled by the parameter n ; at $s = 0$, we recover the original tent map at all n . The physical probability distributions – computed empirically over a trajectory of length 10 billion – at $s = 0.1$, are shown on the right column of Figure 3-20. On the top row, $n = 0$, and there is already a marked asymmetry developed in the physical distribution compared to the uniform distribution, which is the physical distribution seen at $s = 0, n = 0$. The figure shows that by increasing n , despite the fact that

the magnitude of the perturbation becomes smaller with n , we see a dramatic change in the appearance of the stationary probability distribution: we observe an apparent fractal distribution reminiscent of that of the quasi-physical solutions in section 3.4.3. By construction of the map (Supplementary Material section 3.6.3), the nonuniformities of the probability distribution at $n = 0$ are transferred to smaller and smaller scales, as n is increased. It is worth emphasizing that these remarkable changes in the physical distribution (shown on the bottom-right of Figure 3-20) are effected with a tiny perturbation – the perturbed map at $n = 6$ and the original tent map appear indiscernible on the bottom-left of Figure 3-20.

In view of the plucked tent map, we can further question the validity of numerical solutions. Can numerical solutions play the role of the plucking perturbation to the governing equation? That is, can numerical solutions represent slight perturbations to the governing equations for which the physical distribution is drastically different from the physical distribution of the original governing equation? Sauer [240] has shown examples in which numerical error due to computation in double-precision floating point arithmetic causes significant change in the stationary probability distribution on the attractor; other works [170] further analyze the unreliability due to double-precision arithmetic of statistical measures such as power spectra and correlation functions. The plucked tent map adds to this list of examples by demonstrating a mechanism to produce an extreme non-smooth response (Supplementary Material section 3.6.3). Uniformly hyperbolic dynamical systems exhibit *linear response* ([232][28]) by which small parameter perturbations lead to small changes in statistics, which can be expanded as Taylor series around the reference parameter value. But, uniform hyperbolicity is a mathematical idealization, and although some physical systems have been observed to behave as if they were uniformly hyperbolic [108], a violation of uniform hyperbolicity is more likely [275][191].

Both the nonphysicality of shadowing solutions, and the existence of extremely non-smooth statistical response, undermine the validity of using shadowing for sensitivity analysis of statistics [273][206][165]. When the goal is to compute derivatives of ensemble averages, where the ensemble is distributed according to the physical mea-

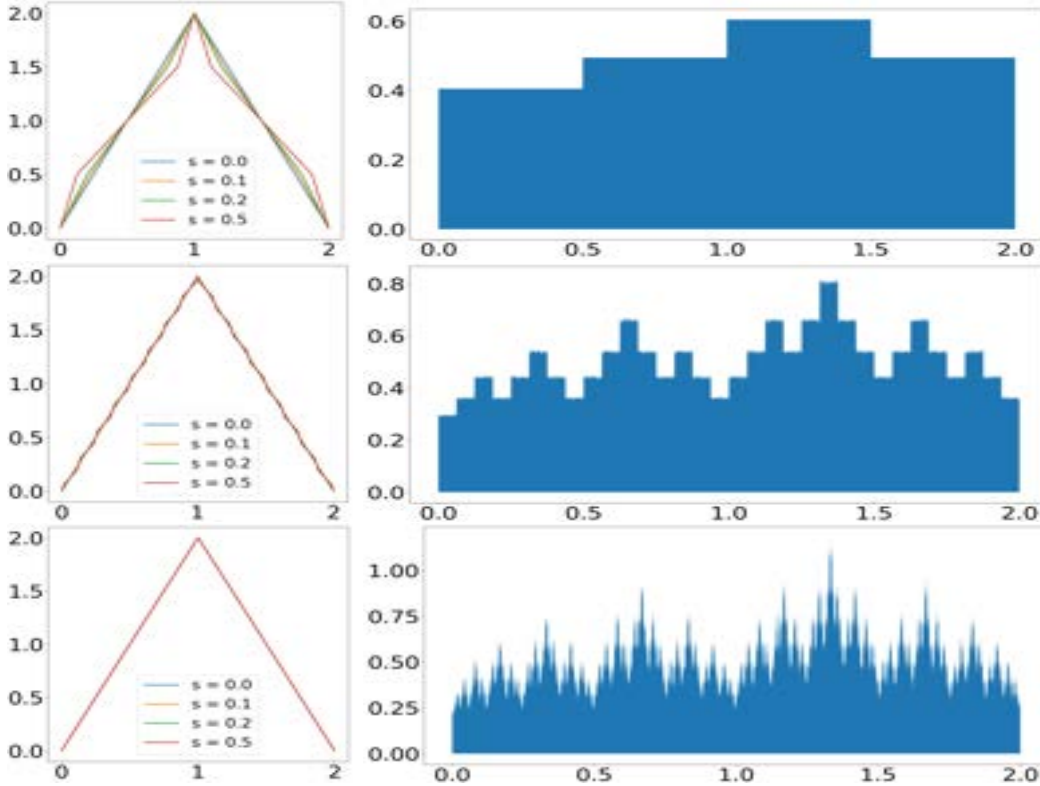


Figure 3-20: The effect of increasing n on the physical probability distribution associated to the plucked tent map. The left column shows the plucked tent map at different values of s and $n = 0$ (top), $n = 3$ (middle), $n = 6$ (bottom). The original tent map is at $s = 0$ on each plot. The right column shows the stationary, physical probability distribution of the plucked tent map at $s = 0.1$ and $n = 0$ (top), $n = 3$ (middle), $n = 6$ (bottom).

sure, i.e., the SRB measure, shadowing-based methods can give wrong results. This is because shadowing-based methods compute the sensitivities of ensemble statistics along shadowing solutions, but these may not be physical solutions that reproduce ensemble statistics. The error in shadowing sensitivities has been observed before and appropriately attributed as “ergodicity breaking error” [37] [209]. Violations of smooth response, like in the plucked tent map, have an immediate implication for shadowing sensitivities as well. In particular, if perturbed solutions are shadowed by physical solutions for all time, the change in the statistics of the perturbed solutions must be small. In other words, the physicality of shadowing predicts that there cannot be a large change in statistics due to small parameter perturbations; the reality is that, as illustrated by the plucked tent map, the effect of small parameter changes

on the statistics can be drastic.

In this chapter, we have constructed several counterexamples that dispel the notion that shadowing solutions are physical solutions, i.e., that their statistical distribution is the same as that typically observed for almost every solution of a governing equation. The existence of long-time shadowing solutions [119] has historically been used to address the issue of whether numerical simulations, which are perturbed solutions, represent the true physics implied by the governing equation. In light of the evidence in this chapter, we must reopen this issue. Even when numerical simulations are shadowed by a true solution in that the difference between them is small for a long time, this shadowing solution may not represent the long-term or ensemble behavior of the physical system. The nonphysicality of shadowing solutions also indicates that shadowing-based methods can lead to incorrect values of sensitivities of statistics to parameter changes.

3.6 Supplementary material

The supplementary material for this chapter including the code and data to generate the figures can be found on Github [56]. The code can be found under the `code` subdirectory inside which section-wise code is separated into further subdirectories. The data used to plot the figures can be found under `data`. The files referred to in this section can be found in the appropriate subdirectory under `code`.

3.6.1 Approximation of the Lorenz Map

The motivation for approximating the Lorenz map is that a closed form expression for the map is necessary for our numerical shadowing procedure (section 3.6.2). In a small region around the cusp, we approximate the map using an exponential function. The tails on both sides are fitted with a sum of a cubic polynomial and a rational function.




Figure 3-21: L: the blue and orange lines are the Lorenz map at the standard parameters, and the dotted black lines indicate the approximate Lorenz map (Table 3.2). R: the regression error is shown as a function of z

Thus, the approximate Lorenz map has the following closed form expression:

$$\varphi(x) = \begin{cases} z_{\max} - f_R(x - z_{\text{sep}}) & x > z_{\text{sep}} \\ z_{\max} - f_L(z_{\text{sep}} - x) & x \leq z_{\text{sep}} \end{cases} \quad (3.21)$$

where

$$f_R(y) = \left((1000y)^p + \sum_{n=0}^3 p_{R,n} y^n + \frac{a_{R,0} + a_{R,1}y}{\sum_{n=0}^3 b_{R,n} y^n} \right), \quad (3.22)$$

and,

$$f_L(y) = \left((1000y)^p + \sum_{n=0}^3 p_{L,n} y^n + \frac{a_{L,0} + a_{L,1}y}{\sum_{n=0}^3 b_{L,n} y^n} \right). \quad (3.23)$$

The location of the cusp is denoted z_{sep} (in which “sep” stands for *separation*), and the maximum and minimum values encountered in the Lorenz map iterates are denoted z_{\max} and z_{\min} , respectively. The exponent of the cusp is denoted p . The coefficients of the cubic polynomial modelling the left (right) tail are denoted $p_{L,n}$, $n = 0, 1, 2, 3$ ($p_{R,n}$, $n = 0, 1, 2, 3$). The coefficients of the numerator and denominator of the rational

function modelling the left (right) tail are denoted $a_{L,0}, a_{L,1}$ ($a_{R,0}, a_{R,1}$) and $b_{L,n}, n = 0, 1, 2, 3$ ($b_{R,n}, n = 0, 1, 2, 3$), respectively. The values of these coefficients, which are obtained by regression, are shown for the map at standard parameters, in the table below. As shown in Figure 3-21, the fit obtained matches the Lorenz map closely.

$z_{\text{sep}} = 38.55302437476555$	$z_{\text{min}} = 29.213182255013322$
$z_{\text{max}} = 47.978140718671284$	$p = 0.28796740575434676$
$p_{L,3} = -0.00024683786275242047$	$p_{L,2} = 0.016174566354858824$
$p_{L,1} = 0.40179772568004946$	$p_{L,0} = -0.24612651488351725$
$p_{R,3} = -0.00020712463321688308$	$p_{R,2} = 0.017130843276711716$
$p_{R,1} = 0.3930080703420676$	$p_{R,0} = -0.23471384266765036$
$a_{L,1} = -0.05405742075580959$	$a_{L,0} = -0.05405742075580959$
$a_{R,1} = -0.05351127783496397$	$a_{R,0} = 0.22489891059122896$
$b_{L,3} = 0.5609397451213353$	$b_{L,2} = -0.3491184293228338$
$b_{L,1} = 2.419972619058592$	$b_{L,0} = 1.0$
$b_{R,3} = 0.6456076059873844$	$b_{R,2} = -0.34840383986411055$
$b_{R,1} = 2.6035438510917692$	$b_{R,0} = 1.0$

Table 3.2: The fitting parameters of the Lorenz map at $\sigma = 10, \beta = 8/3$, and $\rho = 28$.

3.6.2 Computing shadowing solutions

In general, to numerically compute the shadowing solutions, one could use existing methods such as the least squares shadowing method [273]. However, since the maps we consider are all one-dimensional chaotic systems that have in common non-invertibility with two inverse *branches*, we devise a simpler approach. The map φ_s in this section refers to any of the perturbations of the tent map (section 3.4.3) or the Lorenz map. Suppose we are given a perturbed solution $x_n, n = 0, 1, \dots, N$ that we must compute a shadow of. The shadowing solution, $y_n, n = 0, 1, 2, \dots, N$ must a) satisfy the governing equation: $y_{n+1} = \varphi_s(y_n)$ and b) lie close to the given perturbed solution at all times up to N , i.e., $|x_n - y_n| < \epsilon$, for some $\epsilon > 0$, for all $n \leq N$. Suppose we set $y_0 = x_0$, the difference $y_1 - x_1$ will amplify on further iterations under

φ_s for any perturbed solution. On the other hand, backward iteration of the map is contracting, and thus we set $y_N = x_N$. Then, we proceed backward in time to construct a solution of φ_s , noting that any difference that emerges at a given time n will be made smaller starting at $n - 1$. Each point has two pre-images under φ_s , and thus there are two possible choices for y_{N-1} , each lying in one of two fixed sub-intervals separated by the cusp of φ_s . Due to the contraction of errors backward in time, as long as we choose a pre-image in the same subinterval as the perturbed solution, we are guaranteed to approximate a shadowing solution. Marching backward by choosing at each step, the pre-image y_n in the same subinterval as x_n , y_n approximates the shadowing solution better as n decreases. Thus, the procedure to find a shadowing solution simply reduces to solving for a backward trajectory (specifically one among the possible 2^N), starting at a given final condition, x_N .

Hence, it is clear that all we need is the inverse of φ_s , which is propagated backward by choosing the same *branch* of the inverse as x_n at time n . This logic is implemented for each map of section 3.4.3 in the function `shadow` that can be found in the files named for each map (for example, the shadowing solution of the pinched tent map can be found by executing the `shadow` function of `tent_shadow/tent_shadow_pinched.py`). These functions use the analytical inverses of the maps, which are easy to derive for the tent map perturbations of section 3.4.3. For the Lorenz map, we use Newton's method to solve for the inverse, and this is implemented in the file `lorenz_map/shadow.py`. Note that we need a closed form expression of the map, for the Newton's method, and this is indeed the reason why we approximate the map, as described in section 3.6.1.

3.6.3 The *plucked* tent map

We provide a recursive definition of the plucked tent map, which is illustrated at different values of n and s in Figure 3-20. First we define a function $f_s(x)$, which

creates a bend in the tent map that increases with s , around $((1 - s)/2, 1)$:

$$f_s(x) = \min\left(\frac{2x}{1-s}, 2 - \frac{2(1-x)}{1+s}\right), \quad x < 1. \quad (3.24)$$

Then, we introduce oscillations using repetitions of the above map within the unit interval,

$$o_s(x) = \begin{cases} f_s(2x)/2, & x < 0.5 \\ 2 - f_s(2 - 2x)/2, & x \geq 0.5. \end{cases} \quad (3.25)$$

We can modulate the frequency of repetitions – proportional to n – of the oscillations through

$$\lambda_{s,n}(x) = \frac{o_s(2^n x - \lfloor 2^n x \rfloor)}{2^n} + 2 \frac{\lfloor 2^n x \rfloor}{2^n}, \quad (3.26)$$

where $\lfloor x \rfloor$ is the greatest integer less than or equal to x . Finally, the plucked tent map is defined as the above function in the unit interval, and as its reflection about $x = 1$, in the interval $[1, 2)$.

$$\varphi_{s,n}(x) = \min(\lambda_{s,n}(x), \lambda_{s,n}(2 - x)), \quad 0 < x < 2. \quad (3.27)$$

Recall that our motivation is to construction a slight perturbation of the tent map whose stationary probability distribution is not just non-uniform (e.g. like the other tent map perturbations of section 3.4.3), but in which the nonuniformity can be controlled, and made skewed as desired. Thus, we choose to construct a base perturbation of the tent map, $\min(o_s(x), o_s(2 - x))$ – which is also obtained by setting $n = 0$ in $\varphi_{s,n}$ – in which an asymmetry is produced in the probability distribution (top-right of Figure 3-20). Upon repeating, with appropriate scaling, the oscillation o_s , the frequency of which is controlled by n , we obtain $\varphi_{s,n}$, as indicated by Eq. 3.26-3.27. Thus, the asymmetric probability distribution, through this process of scaled repetition, acquires an apparent fractal-like structure seen on increasing n , as shown

on the bottom-right of Figure 3-20.

One way to see the asymmetry in the probability distribution about $x = 1$, at $n = 0$, is to construct a Markov chain with nodes 00, 01, 10, and 11, that indicate the subintervals $(0, 0.5)$, $[0.5, 1)$, $[1, 1.5)$ and $[1.5, 2)$, respectively. From Eq. 3.27, we can see that the application of $\varphi_{s,0}$ yields the following transition matrix:

$$\begin{bmatrix} (1-s)/2 & (1+s)/2 & 0 & 0 \\ 0 & 0 & (1+s)/2 & (1-s)/2 \\ 0 & 0 & (1+s)/2 & (1-s)/2 \\ (1-s)/2 & (1+s)/2 & 0 & 0 \end{bmatrix}.$$

The stationary probability distribution of this Markov chain, which is the left eigenvector of the transition matrix corresponding to eigenvalue 1, is $[(1-s)^2/2, (1-s^2)/2, (1+s)^2/2, (1-s^2)/2]^T$. One can verify this is consistent with the top-right of Figure 3-20 for $\varphi_{0.1,0}$. Clearly, the time spent by an infinitely long physical solution in the second quarter interval, $[0.5, 1]$ is about 1.2 times the time spent in the first quarter, $[0, 0.5]$; the time spent in the third quarter $[1, 1.5]$ is about 1.5 times that spent in the first quarter. This nonuniformity is more pronounced with increasing s . For instance, at $s = 0.5$, the left subinterval, $[0, 1)$, is three times less likely to be visited by a long trajectory compared to $[1, 2)$.

But, as mentioned in the main text, even a small perturbation – very small s – is sufficient to trigger a remarkable variation in the statistics. This is because this particular construction transfers the nonuniformities in the probability distribution at larger scales to smaller scales, with increasing n . The nonuniformity at the largest scale of half intervals is retained at higher values of n . That is, the probability of visiting the interval $[0, 1]$ is the same: $(1-s)$, at all n . However, with increasing n , the nonuniformity emerges within smaller subintervals, due to the construction of the map that relies on repeating the behavior at larger scales at smaller scales (Eq. 3.26). This repetitive construction is responsible for the apparent fractal structure of the probability distribution (Figure 3-20, bottom-right). The scripts that generate the plucked tent map and its physical probability distributions can be found under

tent_sens_stat.

Chapter 4

Ruelle's formula and the infeasibility of its direct evaluation

In chaotic systems, such as turbulent flows, the solutions to tangent and adjoint equations exhibit an unbounded growth in their norms. This behavior renders the instantaneous tangent and adjoint solutions unusable for sensitivity analysis. The Lea-Allen-Haine ensemble sensitivity (ES) estimates provide a way of computing meaningful sensitivities in chaotic systems by utilizing tangent/adjoint solutions over short trajectories. In this chapter, we analyze the feasibility of ES computations under optimistic mathematical assumptions on the flow dynamics. Furthermore, we estimate upper bounds on the rate of convergence of the ES method in numerical simulations of turbulent flow. Even at the optimistic upper bound, the ES method is computationally intractable in each of the numerical examples considered. This chapter is joint work with Pablo Fernandez, Chaitanya Talnikar and Qiqi Wang and has been published at [69].

4.1 Introduction

Gradient-based computational approaches in multi-disciplinary design and optimization require sensitivity information computed from numerical simulations of fluid flow. In Reynolds-averaged-Navier-Stokes (RANS) simulations, sensitivity computation is

traditionally performed using tangent or adjoint equations or using finite difference methods. Sensitivities computed from RANS simulations and non-chaotic Navier-Stokes simulations have been extensively applied toward uncertainty quantification [216, 272], mesh adaptation [104], flow control [229], noise reduction [43, 93, 263] and aerostructural design optimization applications [219, 111, 112, 224, 212]. Many modern applications require computing sensitivities in direct numerical simulations (DNS) or large-eddy simulations (LES); examples include buffet prediction in high-maneuverability aircraft, modern turbomachinery design and jet engine and airframe noise control. Conventional tangent/adjoint approaches cannot be used to compute sensitivities of statistical averages (or long-time averaged quantities) in these high-fidelity, eddy-resolving simulations. This is because the tangent and adjoint solutions diverge exponentially [206, 168], since these simulations exhibit chaotic behavior i.e., infinitesimal perturbations to initial conditions grow exponentially in time. For this reason, sensitivity studies on DNS or LES have been restricted to short-time horizons; these short-time sensitivities have found limited applicability including in flow control in combustion systems [52], jet noise reduction [43] and structural design [201].

The first proposed approach to tackle the computation of sensitivities of statistics is the Lea-Allen-Haine ensemble sensitivity (ES) method [168]. In this method, the problem of exponentially diverging linearized perturbation solutions (such as tangent/adjoint) is mitigated by taking a sample average of short-time sensitivities. The method approximates Ruelle’s response formula [232] for sensitivity of average quantities to system parameters. The convergence of the method has been shown by Eyink *et al.* [96] in the limit of taking an infinite number of samples and increasing the time duration for sensitivity computation to infinity, in that order. Eyink *et al.* [96] establish that the rate of convergence is worse than a typical Monte Carlo simulation (in which the error in a sample average diminishes at the rate $1/\sqrt{N}$, where N is the number of samples) in the case of the classical 3-variable Lorenz ’63 system. However, the convergence trend is still unknown for general chaotic systems. In this work, we present an analysis of the mean squared error of the ES method as a function of the computational cost for a certain class of systems called uniformly hyperbolic systems

[154]. It is worth noting that at the time of writing of this chapter, alternatives to the ES method [210, 67, 211] are under active investigation. Non-intrusive least squares shadowing (NILSS) [206, 37] and its adjoint-variant [211, 37] are methods that are conceptually based on the shadowing property of uniformly hyperbolic systems. This property enables the computation of a particular tangent solution that remains bounded in a long time window and sensitivities are estimated using this tangent solution. The NILSS algorithm requires the knowledge of the unstable subspace (of the tangent space, corresponding to the positive Lyapunov exponents). This makes the algorithm expensive when the dimension of the subspace or the number of positive Lyapunov exponents is large. The NILSS method has been applied to LES of turbulent flows around bluff bodies [39, 40] at low Reynolds numbers, where the number of positive Lyapunov exponents is small enough to limit the computational expense when compared to wall-bounded flows, for instance.

The chapter is organized as follows. In the next section, we review the ES method and define the ES estimator for the sensitivity. In section 4.3, we describe the mean squared error in the ES estimator in terms of the associated bias and variance and obtain optimistic, problem-dependent estimates for both components. We predict the rate of convergence as a function of computational cost under these optimistic assumptions on the dynamics. The rest of the chapter consists of numerical examples that illustrate the convergence trend of the ES method. In sections 4.4.1 and 4.4.2, we discuss two low-dimensional models of chaotic fluid behavior: the Lorenz '63 and Lorenz '96 systems. We apply our optimistic analysis to roughly estimate an upper bound on the rate of convergence. We then present two numerical simulation results that serve to illustrate the applicability of ES schemes in fluid simulations of practical interest, in light of our mathematical analysis in section 4.3. The first is a simulation of a NACA 0012 airfoil in section 4.4.3 and the second, an LES of turbulent flow around a turbine vane in section 4.4.4. Section 4.5 contains some practical recommendations on the applicability of the ES method, based on analytical and numerical insights from the previous sections.

4.2 The ES estimator

4.2.1 The sensitivity computation problem setup

Consider a chaotic fluid flow parameterized by s , expressed through a PDE system of the following form:

$$\begin{aligned} \frac{\partial u(t, s)}{\partial t} &= f(u(t, s), s) & u(t, s) \in \mathbb{R}^d, s \in \mathbb{R}^p, \\ u(0, s) &= u_0 \in \mathbb{R}^d. \end{aligned} \tag{4.1}$$

Here, we use $u(t, s)$ to denote the state vector at time t obtained due to the evolution of an initial state u_0 according to Eq. 4.1; u_0 is chosen independently of s . As an example, if Eq.4.1 is the spatially discretized incompressible Navier-Stokes equation, a state vector consists of the velocity components, and pressure at all the grid points. In this case, the dimension of the parameterized, spatially-discretized Navier-Stokes state vector – $u(t, s)$ – is $d = 4 \times$ the number of degrees of freedom (and in a two-dimensional flow, $d = 3 \times$ the degrees of freedom). The state vector $u(t, s)$ is approximately known from numerical simulation and can be thought of as a point in the phase space M , a compact subset of \mathbb{R}^d . The set of input parameters s can be, for instance, related to the inlet conditions, the geometry of the domain or solid bodies in the flow and so on. In the interest of simplicity, from here on, s is a scalar parameter. The state vector is also a function of the initial condition and to make explicit this dependence, we write as $u(t, s, u_0)$ the solution of Eq.4.1 at time t , solved with u_0 as the initial state.

The fluid flows we consider here are statistically stationary, i.e., the states in phase space are distributed according to a time-invariant probability distribution μ_s . The subscript s indicates that the stationary distribution is a function of the parameter; we work under the assumption that μ_s is a smooth function of s . Suppose J is a smooth scalar function of the state such as the lift/drag ratio or the pressure loss in a turbine wake. The ensemble mean of J is defined as its expectation with respect to μ_s and denoted by $\mu_s(J) := \int J d\mu_s$. We are interested in computing the sensitivity of

$\mu_s(J)$ to s . The ensemble mean $\mu_s(J)$, is an average over phase space, and hence can be measured as a time average along almost every trajectory, under the assumption of ergodicity. More precisely, for almost every initial state u_0 ,

$$\mu_s(J) = \lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t J(u(t', s, u_0)) dt'.$$

In practice, $\mu_s(J)$ is computed approximately as a finite-time average by truncating a trajectory at a large t . The Lea-Allen-Haine ES method, the subject of this chapter, computes the sensitivity $(d\mu_s(J)/ds)$ approximately, as we describe next.

4.2.2 The Lea-Allen-Haine ES algorithm

The ES estimator of the sensitivity is the sample mean of a finite number of independent sensitivity outputs computed over short trajectories. We now make this statement precise in the following description of the ES algorithm. Consider N independent initial states $\{u_0^{(i)}\}_{i=1}^N$, sampled according to μ_s . Let us denote the sensitivity computed along a flow trajectory, of length τ , starting from $u_0^{(i)}$, as $\theta_\tau^{(i)}$. Then, the Lea-Allen-Haine estimator $\theta_{\tau,N}$, is given by,

$$\theta_{\tau,N} = \frac{1}{N} \sum_{i=1}^N \theta_\tau^{(i)}. \quad (4.2)$$

The standard adjoint method was proposed to be used originally [168, 96] to compute the sensitivities $\theta_\tau^{(i)}$ in order to retain the advantage of adjoint methods, namely that they scale well with the parameter space dimension, since the parameters enter the computation only when determining the sensitivities using the adjoint solution vectors, and the adjoint vectors the computation of which is the majority of the computational cost, are themselves parameter-independent. The analysis in the remainder of this chapter would be identical however, if the tangent equation or a finite difference approximation to the sensitivity derivative was used instead. The three methods of computing $\theta_\tau^{(i)}$, dropping the superscript i for clarity, are listed below. The sensitivity estimator computed using Eq. 4.2 when $\theta_\tau^{(i)}$ are computed using the

tangent equation, adjoint equation and from finite difference are denoted using $\theta_{\tau,N}^T$, $\theta_{\tau,N}^A$ and $\theta_{\tau,N}^{FD}$ respectively.

1. From the tangent equation:

$$\frac{dv(t, u_0)}{dt} = \frac{\partial f}{\partial s} \Big|_t + \frac{\partial f}{\partial u} \Big|_t v(t, u_0) \quad (4.3)$$

where, $v(t, u_0) = (du(t, s, u_0)/ds)$ is the tangent solution at time t when the primal initial condition is u_0 . The tangent initial condition $v(0, u_0) = 0 \in \mathbb{R}^d$ since u_0 is independent of s . We use $(\partial f/\partial s) \Big|_t$ to represent the partial derivative of f with respect to the parameter s , evaluated at $u(t, s, u_0)$. The Jacobian matrix at time t is written as $(\partial f/\partial u) \Big|_t$. The ES estimator based on the tangent equation is

$$\theta_{\tau,N}^T = \frac{1}{\tau N} \sum_{i=1}^N \int_0^\tau \frac{\partial J}{\partial u} \Big|_{t,i} \cdot v(t, u_0^{(i)}) dt, \quad (4.4)$$

where we have used a second subscript i in $\frac{\partial J}{\partial u} \Big|_{t,i}$ to indicate that the primal initial condition is $u_0^{(i)}$.

2. From the adjoint equation:

$$\frac{dy(t, u_0)}{dt} = -\frac{\partial J}{\partial u} \Big|_t - \left(\frac{\partial f}{\partial u} \right)^* \Big|_t y(t, u_0), \quad (4.5)$$

where $y(\tau, u_0) = 0 \in \mathbb{R}^d$ is the adjoint vector at time τ which results in the adjoint vector $y(t, u_0)$ at time t when evolved backwards in time for a time $\tau - t$. Here $(\partial f/\partial u)^* \Big|_t$ is the conjugate transpose of the Jacobian matrix at time t . The ES estimator based on the adjoint equation is,

$$\theta_{\tau,N}^A = \frac{1}{\tau N} \sum_{i=1}^N \int_0^\tau y(t, u_0^{(i)}) \cdot \frac{\partial f}{\partial s} \Big|_{t,i} dt, \quad (4.6)$$

where we have used a second subscript i in $(\partial f/\partial s) \Big|_{t,i}$ to indicate that the

primal initial condition is $u_0^{(i)}$.

3. Using, for example, a second-order accurate finite difference approximation:

$$\theta_{\tau,N}^{\text{FD}} = \frac{1}{\tau N} \sum_{i=1}^N \frac{1}{2\Delta s} \left(\int_0^\tau J(u(t, s + \Delta s, u_0^{(i)})) dt - \int_0^\tau J(u(t, s - \Delta s, u_0^{(i)})) dt \right), \quad (4.7)$$

for a small Δs .

4.3 Error vs. computational cost of the ES estimator

Chaotic systems such as turbulent fluid flows often exhibit regularity in long-time averages despite showing seeming randomness in instantaneous measurements. The chaotic hypothesis, proposed by Gallavotti and Cohen [108], is the notion that these systems can be treated, for the purpose of studying their long-time behavior, as having a certain smooth structure in phase space. This smooth structure allows for the existence of subspaces of the tangent space consisting of expanding and contracting derivatives of state functions. Our goal in this section is to predict the convergence trend of the ES method for systems that satisfy the chaotic hypothesis. More specifically, we would like to construct an optimistic model for the least mean squared error of the ES method, achievable at a given computational cost, in these systems. Before we delve into the construction of the optimistic model, we will discuss the rigorous justification for the convergence of the ES method – Ruelle’s response formula. Here we will focus on the implications of the formula for the ES method without going into details; the reader is referred to Ruelle’s original chapter [232, 234] for the derivation of the response formula. We shall refrain from a technical discussion on hyperbolicity and other concepts from dynamical systems theory but provide the necessary qualitative description, in the context used here.

4.3.1 ES estimator as approximation of Ruelle's response formula

The following response formula [235] due to Ruelle gives the sensitivity to s , of an objective function J 's statistical average,

$$\frac{d\mu_s(J)}{ds} = \int_0^\infty dt \int_M \left. \frac{\partial J}{\partial u_0} \right|_t \cdot \left. \frac{\partial f}{\partial s} \right|_0 d\mu_s(u_0), \quad (4.8)$$

where $\left. (\partial J / \partial u_0) \right|_t := \partial J(u(t, s; u_0)) / \partial u_0$ represents the derivative of J at time t with respect to the initial state u_0 . One can interpret the inner integral (over M) as the statistical response of the objective function at time instant t . It is a phase space average of the sensitivity at time t with the initial conditions distributed according to the stationary probability distribution μ_s . Ruelle's formula has been proven to hold for a certain class of smooth dynamical systems known as uniformly hyperbolic systems. Roughly speaking, these are systems in which the tangent space at every point in phase space can be split into stable and unstable subspaces, which contain respectively, exponentially decaying and growing tangent vectors. In reality, the formula is applicable to a large class of fluid flow problems (this wider applicability referenced earlier as the chaotic hypothesis) that are not necessarily uniformly hyperbolic, as subsequent works [108, 232] have analyzed.

An iterated integral such as in Eq. 4.8, gives the same value upon switching the order of integration if and only if the double integral, in which the integrand is replaced by its absolute value, is finite (this is the Fubini-Tonelli theorem). The absolute value of the integrand in Eq. 4.8 diverges to infinity as $t \rightarrow \infty$ for almost every initial condition in phase space. In other words, in a chaotic system, for every initial condition u_0 in phase space except those in a set of μ_s -measure 0,

$$\int_0^\infty \left| \left. \frac{\partial J}{\partial u_0} \right|_t \cdot \left. \frac{\partial f}{\partial s} \right|_0 \right| dt = \infty.$$

For this reason, in Eq.4.8, the integral over phase space and over time do not commute. The iterated integral in Eq.4.8 in which the integration over phase space is

performed first, leads to a bounded value, which is equal to $d\mu_s(J)/ds$. On the other hand, changing the order of integration and integrating over t first, results in infinity.

From here on, we use the following sample average as the definition of the ES estimator

$$\theta_{\tau,N} := \frac{1}{N} \sum_{i=1}^N \int_0^\tau \frac{\partial J}{\partial u_0} \Big|_{t,i} \cdot \frac{\partial f}{\partial s} \Big|_{0,i} dt, \quad (4.9)$$

where the set of initial states $\{u_0^{(i)}\}_{i=1}^N$ is independent and identically distributed according to μ_s . The ES estimator can be interpreted as an approximation of Ruelle's formula in the following sense: if the outer integral over time (in Eq. 4.8) is truncated at time τ and the phase space average of the integrand approximated with a sample mean over N independent samples, we obtain the estimator in Eq. 4.9. The definition can also be interpreted as the average sensitivity of $\int_0^\tau J(u(t, s, u_0)) dt$ to initial condition perturbations along the vector field $(\partial f/\partial s) \Big|_0$. This can be computed using the tangent equation method listed in 4.2.2 but using the homogeneous tangent equation – i.e., the tangent equation without the forcing term, $(\partial f/\partial s) \Big|_t$. To wit, the homogeneous tangent equation solved with the initial condition $v^h(0, u_0) = (\partial f/\partial s) \Big|_0$ yields at time t , the solution,

$$v^h(t, u_0) = \frac{\partial u}{\partial u_0} \Big|_t \frac{\partial f}{\partial s} \Big|_0.$$

Therefore, the sensitivity of $\int_0^\tau J(u(t, s, u_0)) dt$ to initial condition perturbations along $(\partial f/\partial s) \Big|_0$ is,

$$\int_0^\tau \frac{\partial J}{\partial u_0} \Big|_t \cdot \frac{\partial f}{\partial s} \Big|_0 dt = \int_0^\tau \frac{\partial J}{\partial u} \Big|_t \cdot \frac{\partial u}{\partial u_0} \Big|_t \frac{\partial f}{\partial s} dt = \int_0^\tau \frac{\partial J}{\partial u} \Big|_t \cdot v^h(t, u_0).$$

Taking an N -sample average of the above results in the formula for the estimator 4.9. The resulting $\theta_{\tau,N}$ is different from the sensitivities $\theta_{\tau,N}^{\text{FD}}$, $\theta_{\tau,N}^{\text{T}}$ and $\theta_{\tau,N}^{\text{A}}$ defined in section 4.2.2, all of which are also different from one another. However, in the asymptotic limit of $\tau \rightarrow \infty$, all these sensitivities grow exponentially ($\theta_{\tau,N}^{\text{FD}}$ does

not grow unbounded in norm with τ , but rather saturates, for a non-zero value of Δs) at the same rate determined by the largest among the Lyapunov exponents (the asymptotic exponential growth or decay rate of tangent/adjoint vectors) of the system. Therefore, for the purpose of an asymptotic analysis, we restrict our attention to the estimator defined by Eq. 4.9 and refer to $\theta_{\tau,N}$ using the umbrella term ES estimator.

One notices that in the practical computation of the ES method, the integrals are commuted when compared to Ruelle’s formula but the integral in time is truncated at a finite time. As noted in Eyink et al’s analysis [96], the rationale behind swapping the order of integration as compared to Ruelle’s formula in Eq. 4.8 is the observation that the “divergence of the individual adjoints is delayed on taking a sample average”, for the Lorenz ’63 system, a low-order model for fluid convection that we discuss in section 4.4.1. In the rest of the chapter, our goal is to analyze the convergence trend in more generality.

4.3.2 Bias and variance of the ES estimator

Having defined the ES estimator, we now construct an optimistic model for its mean squared error in uniformly hyperbolic systems. In this section, we present our choices of optimistic estimates for the bias and the variance associated with the estimator and use uniform hyperbolicity to justify those choices. The ES estimator, as defined in Eq.4.9, has a non-zero bias for a finite τ . By definition, the bias, denoted by $b(\tau)$, gives the difference between the value attained by the estimator on using an infinite number of samples and the true value of the sensitivity,

$$b(\tau) = \mu_s(\theta_{\tau,N}) - \frac{d\mu_s(J)}{ds}. \quad (4.10)$$

Therefore b is only a function of τ (and not N) because on using an infinite number of samples,

$$\lim_{N \rightarrow \infty} \theta_{\tau,N} = \mu_s(\theta_{\tau,N}) = \int_M d\mu_s \int_0^\tau \frac{\partial J}{\partial u_0} \Big|_t \cdot \frac{\partial f}{\partial s} dt. \quad (4.11)$$

Writing Eq. 4.10 more explicitly as,

$$b(\tau) = \int_M \int_0^\tau \left. \frac{\partial J}{\partial u_0} \right|_t \cdot \left. \frac{\partial f}{\partial s} \right|_t dt d\mu_s - \int_0^\infty \int_M \left. \frac{\partial J}{\partial u_0} \right|_t \cdot \left. \frac{\partial f}{\partial s} \right|_t d\mu_s dt, \quad (4.12)$$

clearly indicates why there is a non-zero bias for a finite value of τ . As an optimistic estimate of $b(\tau)$, we choose an exponential decay at a problem-dependent rate denoted γ_1 , using the following justification. In a uniformly hyperbolic system, a tangent vector can be decomposed, at every point in phase space, into its stable and unstable components. A stable (unstable) tangent vector would diminish in norm along a forward (backward) trajectory at an exponential rate. More precisely, for almost every u_0 , if $v^h(0, u_0)$ is a stable tangent vector at u_0 , there exist $C, \alpha > 0$ such that,

$$\|v^h(t, u_0)\| \leq C \exp(-\alpha t) \|v^h(0, u_0)\|, \quad \text{for all } t \geq 0. \quad (4.13)$$

That is, in uniformly hyperbolic systems, solving the homogeneous tangent equation with a stable tangent vector as the initial condition, results in a stable tangent vector whose norm grows smaller exponentially. It must be mentioned here that the interval $[-\alpha, \alpha]$ (with α defined by Eq. 4.13) lies in the gap between the largest negative and smallest positive Lyapunov exponents. That α provides a lower bound for the smallest positive Lyapunov exponent will be used in our analysis later. Let us now examine the bias when the initial perturbation given by $(\partial f / \partial s)(u_0)$ is a stable tangent vector at every u_0 . From Eq. 4.12,

$$b(\tau) = - \int_M \int_\tau^\infty \left. \frac{\partial J}{\partial u_0} \right|_t \cdot \left. \frac{\partial f}{\partial s} \right|_0 dt d\mu_s = - \int_M \int_\tau^\infty \left. \frac{\partial J}{\partial u} \right|_t \cdot v^h(t, u_0) dt d\mu_s \quad (4.14)$$

$$\implies |b(\tau)| \leq \int_M \int_\tau^\infty \left\| \left. \frac{\partial J}{\partial u} \right|_t \right\| (C \exp(-\alpha t)) \left\| \left. \frac{\partial f}{\partial s} \right|_0 \right\| dt d\mu_s \quad (4.15)$$

$$\leq \frac{C}{\alpha} \left\| \left. \frac{\partial J}{\partial u} \right|_\infty \right\| \left\| \left. \frac{\partial f}{\partial s} \right|_\infty \right\| \exp(-\alpha\tau). \quad (4.16)$$

The norm $\|\cdot\|_\infty$ of a vector field X expressed in coordinates as $X(u) = (X_1(u), X_2(u), \dots, X_d(u))$ is defined as $\|X\|_\infty := \sup_{u \in M} \|X(u)\|$, where the norm of the tangent vector $X(u)$ is

the l^2 -norm, $\|X(u)\| := \sqrt{(\sum_{i=1}^d |X_i(u)|^2)}$. We obtain Eq. 4.14 by recognizing that the order of integration can be swapped in the second term (the true sensitivity) in Eq. 4.12 in this case since the absolute value of the integrand is bounded for all time. Equation 4.15 follows from using the Cauchy-Schwarz inequality and the fact that stable perturbations decay in time, as described by inequality 4.13. Thus, we obtain that the bias of the estimator $\theta_{\tau,N}$ decays exponentially with τ if the perturbations lie entirely in the stable subspace of the tangent space at every point. In general, the initial tangent vector will also have an unstable component. The unstable contribution to the bias follows the decay of time correlations [232] in the system, which has been shown to be exponential, at best, in uniformly hyperbolic systems. Collet and Eckmann [78] proposed the conjecture that for randomly picked observables in expanding systems (where all the Lyapunov exponents are strictly positive), the decay of time correlations is at an exponential rate that is at least as slow as the smallest positive Lyapunov exponent. Applying the Collet-Eckmann conjecture provides another rationalization of the fact that even in the optimistic event of the perturbation field being stable (Eq.4.14 - Eq.4.16), we would expect the time for the decay of the bias (and hence the minimum time required for the convergence of the ES method) to be at least that of the decay of time correlations in the system. This is because the exponential rate at which the bias decays in the stable perturbation case is α (from Eq.4.16), a lower bound on the smallest positive Lyapunov exponent, which assuming the Collet-Eckmann conjecture applies, is on the same order as the decay of time correlations in the system.

Estimating the decay of correlations is an active research area and previous studies [26, 27] have obtained that even among hyperbolic chaotic systems, *intermittent* systems can exhibit subexponential decay of correlations. Therefore, an exponential decay of the bias with integration time is justified as a representation of the optimal scenario, giving rise to the following model for the squared bias, for some constant $C_b > 0$:

$$b^2(\tau) = C_b \exp(-2\gamma_1\tau), \quad (4.17)$$

where γ_1 is a problem-dependent rate. In the same vein as our discussion on the bias above, we propose a model for the best case variance and provide a justification for the chosen model. The ES estimator $\theta_{\tau,N}$ is a sample average of the random variable $\int_0^\tau (\partial J/\partial u_0)|_t \cdot (\partial f/\partial s)|_0 dt$, the randomness arising in the (deterministic) chaotic system due to the randomness in the initial condition. We know that the initial conditions are distributed according to μ_s and this gives rise to an unknown τ -dependent distribution for $\int_0^\tau (\partial J/\partial u_0)|_t \cdot (\partial f/\partial s)|_0 dt$. For a finite τ , we assume that the variance of this distribution is finite. Then, it follows that since $\mu_s(\theta_{\tau,N})$ is bounded as we established above, the central limit theorem (CLT) applies and therefore, for large N ,

$$\text{var}(\theta_{\tau,N}) \rightarrow \frac{\text{var}(\int_0^\tau (\partial J/\partial u_0)|_t \cdot (\partial f/\partial s)|_0 dt)}{N}. \quad (4.18)$$

The applicability of the CLT for the distribution of $\theta_{\tau,N}$ is in general an optimistic assumption, as discussed in previous works [96] and in the numerical example in section 4.4.1. It is reasonable to expect that $\text{var}(\int_0^\tau (\partial J/\partial u_0)|_t \cdot (\partial f/\partial s)|_0 dt)$ increases exponentially with τ since for almost every initial state, $|(\partial J/\partial u_0)|_t \cdot (\partial f/\partial s)|_0| \sim \mathcal{O}(\exp(\lambda_1 t))$, where λ_1 is the largest Lyapunov exponent of the system. Therefore, we expect that the variance grows exponentially at the rate of twice the largest Lyapunov exponent of the system. Thus, we propose the following optimistic model for the variance, for some $C_{\text{var}} > 0$,

$$\text{var}(\tau, N) = \frac{C_{\text{var}} \exp(2\lambda_1 \tau)}{N}. \quad (4.19)$$




Figure 4-1: The mean squared error as a function of integration time τ , for different values of the variables k and r . The optimal value τ^* is marked on each of the plots.

4.3.3 Optimistic convergence estimate for the ES method

Using the optimistic estimates for the bias and the variance described in section 4.3.2, we arrive at the following ansatz for the mean squared error in the ES estimator:

$$\tilde{e}(\tau, T) = b^2(\tau) + \text{var}(\tau, N) \quad (4.20)$$

$$= C_b \exp(-2\gamma_1\tau) + \frac{C_{\text{var}}\tau}{T} \exp(2\lambda_1\tau), \quad (4.21)$$

where we use $T := N\tau$ to denote the computational cost. The relationship between the integration time τ^* that minimizes the mean squared error and the cost T for the model in Eq. 4.21 is:

$$\tau^*(T) = \arg \min_{\tau} \tilde{e}(\tau, T) = -\frac{1}{2\lambda_1} + \frac{\mathcal{W}(c)}{2(\gamma_1 + \lambda_1)} \quad (4.22)$$

where,

$$c = 2 C_b \frac{\gamma_1(\gamma_1 + \lambda_1)T}{C_{\text{var}}\lambda_1} \exp(1 + \gamma_1/\lambda_1)$$

and \mathcal{W} is the Lambert W -function. The above relationship shows that, given constants C_b and C_{var} independent of τ , the optimal trajectory length of each independent sensitivity evaluation, τ^* , varies sub-logarithmically with the cost T .

From Eqs. 4.21 and 4.22, it can be seen that the least mean squared error, denoted

by $\tilde{e}_{\min}(T) := \tilde{e}(\tau^*(T), T)$ can be reduced to a function of the variables, $k := C_b\gamma_1/C_{\text{var}}$ and $r := \gamma_1/\lambda_1$. Figure 4-1 shows the variation of $\tilde{e}(\tau, T)$ with τ at a fixed T , for different values of k and r . It is clear that the mean squared error is high in magnitude when $r < 1$ no matter the trajectory length, τ . As expected for $r > 1$, the optimal trajectory length increases with r and leads to smaller mean squared errors. The mean squared error is large, as expected, both when a) τ is so short that variance is negligible, and the bias is significant and b) τ is large enough that the effect of exponential increase in the variance is clear. Thus we see a roughly parabolic shape centered at the optimal trajectory length for $r = 1$, where the competing timescales are equal. The parabola is asymmetric $r < 1$ because the time constant of the squared bias decay ($1/(2\gamma_1)$) is larger than that for the variance increase ($1/(2\lambda_1)$). Thus, the bias drops more slowly when compared to $r \geq 1$, and upon increasing τ , is overwhelmed very quickly by the dramatic increase in the variance, which is reflected in $\tilde{e}(\tau, T)$ for $\tau > \tau^*$. Analogously, the downward deflection of the right arm of the parabola for $r > 1$, is explained by the fact that for values of $\tau > \tau^*$, we have a negligible bias and the exponential increase of the variance is slow, when compared to at $r = 1$, while for $\tau < \tau^*$, increasing the trajectory length dramatically drops the bias but the variance remains small, and hence presents as a quick decline in the mean squared error. In 4-2 (left), $\tilde{e}_{\min}(T)/k$ is shown against the total computational cost T at different values of r . The plots in 4-2 (left) reveal that \tilde{e}_{\min} is shifted upward on decreasing k but the slope of \tilde{e}_{\min} vs. T , is relatively unaffected by a change in k when compared to a change in r . From 4-2 (left), we also observe that the least mean squared error \tilde{e}_{\min} decays like an approximate power law of the total cost T . That is,

$$\tilde{e}_{\min} \sim \mathcal{O}(T^{-\beta}) \quad \text{for some } \beta \equiv \beta(k, r) > 0. \quad (4.23)$$

In 4-2 (right), the rate of convergence β is reported at different values of r and k . From 4-2 (right), β appears to be quite robust to varying the ratio of the bias to variance coefficients, C_b/C_{var} , when γ_1 is kept constant. On the other hand, it can be seen that the influence of the ratio of timescales r is significant on the rate




Figure 4-2: Left: \tilde{e}_{\min}/k as a function of T , at different values of r . Right: Rate of convergence as a function of the ratio $r = \gamma_1/\lambda_1$.

of convergence. For values of $r > 1$, the least mean squared error falls faster than $1/\sqrt{T}$, for a range of values of k , indicating that convergence rates better than a typical Monte Carlo sampling can be achieved on choosing an optimal τ under the assumptions of section 4.3.2. This implies that, assuming a strongly chaotic system satisfying our optimistic estimates, when the timescale of the decay of bias is shorter than that of the growth of perturbations ($1/\lambda_1$), the ES method can be very efficient. On the contrary, 4-2 also indicates that the number of samples required to half \tilde{e}_{\min} at $r = 1$ must be increased four-fold while at $r = 0.1$, a factor of 2^{10} increase in the number of samples is required to half the error. To summarize, even in the ideal case of exponential decay of the bias, a rate of decay smaller than the leading Lyapunov exponent, would lead to a significantly less efficient method than a typical Monte Carlo.

4.4 Numerical examples

The previous section was dedicated to a mathematical analysis of the convergence of the ES method. We were able to predict the best possible rates of convergence

under suitable assumptions on the dynamics. In this section, we treat numerical examples of low-dimensional chaotic systems as well as simulations of turbulent flow. In each of the examples, our goal is to gauge, based on our numerical results, the applicability of the assumptions adopted in our analysis in section 4.3. When deemed applicable, we estimate the rate of convergence using our results from section 4.3 and discuss the computational tractability of the ES method given this rate. In other cases where either the analysis is inapplicable or the estimation of bias and variance is not practical, we use a physics-informed approach to predict the rate of convergence. The discussions following the numerical results delineate guidelines for determining the practicality of the ES method.

4.4.1 The Lorenz '63 attractor

As noted in the introduction, Eyink *et al.* [96] have performed a numerical analysis of the ensemble adjoint and related methods on the Lorenz '63 system and make several important observations regarding the convergence trends of $\theta_{\tau,N}$ in τ and in N . We choose the Lorenz '63 system as the first example in order to validate our present results against Eyink *et al.*'s. The Lorenz '63 system is a 3-variable model of fluid convection that is used as a classic example of chaos [179]. It consists of the following system of ODEs:

$$\begin{aligned}\frac{dx}{dt} &= -\sigma x + \sigma y \\ \frac{dy}{dt} &= -xz + sx - y \\ \frac{dz}{dt} &= xy - bz,\end{aligned}\tag{4.24}$$

with the standard values of $\sigma = 10$, $b = 8/3$ and $s = 28$. These equations were derived by Lorenz and Saltzman [179] from the conservation equations for a fluid between horizontal plates maintained at a constant temperature difference. The phase vector $u := [x, y, z]^T$, whose evolution these equations describe, corresponds to coefficients in the Fourier series expansion of the stream function and the temperature profile. The

parameter s is the Rayleigh coefficient normalized by the critical value above which flow instabilities develop. The objective function of our interest is $J(u(t)) := z(t)$ which is proportional to the deviation of the temperature from the linear profile that would be achieved if the fluid were static. It is well-known that a compact attractor exists in phase space.

We use the algorithms described in section 4.2.2 to compute all three types of ES estimators $\theta_{\tau,N}^{\text{FD}}$, $\theta_{\tau,N}^{\text{A}}$ and $\theta_{\tau,N}^{\text{T}}$. To ensure that the initial condition is sampled from the steady-state distribution on the Lorenz attractor, the system is evolved for a spin-up time of about 1.1 time units, before we start computing the sensitivities. This spin-up time is estimated as $1/\lambda_1$, with $\lambda_1 \approx 0.9$ known from the literature to be the largest Lyapunov exponent. The true value of the sensitivity $(d\mu_s(J)/ds)$ was computed by Eyink *et al.* [96] to be ≈ 0.96 . This value is obtained by numerically computing $\mu_s(J)$ as an ergodic average at a range of values of s . It can be seen that $\mu_s(J)$ turns out approximately to be a straight line with a slope of 0.96 around $s = 28$. The primal, the tangent and adjoint dynamics in equations 3.3, 4.3 and 4.5 respectively are computed using forward Euler time integration with a timestep of 0.005 time units. The computational cost $T = N\tau$ was chosen to be 5000 time units. The estimators were computed for a range of values of τ up to 3 time units. In order to apply our analysis in section 4.3, we wish to numerically estimate the bias and variance of $\theta_{\tau,N}$. We estimate $\mu_s(\theta_{\tau,N})$ as a sample average of $\theta_{\tau,N}$ using 5 million independent samples. That is, we approximate $\mu_s(\theta_{\tau,N})$ as $\theta_{\tau,5 \times 10^6}$ while the true value is not a function of the number of samples but only of τ . The variance of $\theta_{\tau,N}$ is again approximated as a sample average wherein $\mu_s(\theta_{\tau,N})$ is replaced with its estimate $\theta_{\tau,5 \times 10^6}$.

Empirically determined upper bound for rate of convergence

Our overarching goal is to predict the rate of convergence of the ES method. We now discuss that it is possible to roughly estimate the rate based on numerical results at a fixed computational cost T . In Lea-Allen-Haine's application of the ES method on the Lorenz '63 system [168], $\tau = 1$ is used to obtain an accurate estimate of $d\mu_s(J)/ds$;

for the sake of comparison with their results, we estimate the rate of convergence locally around $\tau = 1$. To obtain the rate of convergence, we obtain estimates, for the exponential rates of increase and decrease of the variance and the bias respectively, from our numerical results. Then, we use 4-2 to obtain the rate of convergence from the ratio of these obtained rates.

First we establish an empirical estimate for the variance rate. Our numerical estimates for the bias and variance of the estimators $\theta_{\tau,N}^A, \theta_{\tau,N}^T$ and $\theta_{\tau,N}^{FD}$ computed as described above, are shown in 4-3 (left). In 4-3 (right), the numerical estimates of $\mu_s(\theta_{\tau,N}^2)$ are shown for all three estimators. Ignoring transient behavior for τ up to ~ 0.5 , the asymptotic exponential rate of the variance estimate, obtained by determining the slope from 4-3 (left), is larger than the $2\lambda_1$ rate that we predicted in 4.3.3. To understand this result, let us assert that our numerical estimate of the variance reflects the trend of the true variance and the CLT holds for the given range of values of τ . If the first assertion holds, the true variance is increasing exponentially at a faster rate than $2\lambda_1$. Now, owing to the convergence of Ruelle's formula, we know that $\mu_s(\theta_{\tau,N})$ cannot have an unbounded growth as a function of τ and therefore, the rate of increase of $\text{var}(\theta_{\tau,N})$ must be captured by that of $\mu_s(\theta_{\tau,N}^2)$. Computing the slope from 4-3 (right), we see that the numerical estimate of $\mu_s(\theta_{\tau,N}^2)$ is indeed exponential in τ at the rate $\sim 2 \times 0.85$ for both the tangent and adjoint estimators, from a least-squares fit. This value of the rate is closer to our theoretical prediction of $2\lambda_1$. Therefore, it is reasonable to conclude that neither of our previously stated assertions holds true. As a result, we have considerable error in our estimates of $\mu_s(\theta_{\tau,N})$ and $\mu_s(\theta_{\tau,N}^2)$ since the error is decaying slower than expected from CLT. Since both these errors play a role in the variance estimation, obtaining the rate of increase from the estimate of $\mu_s(\theta_{\tau,N}^2)$ is more accurate. It is thus reasonable to take the better numerical estimate, 2×0.85 , as the rate of exponential increase of the variance for τ up to 1.5.

In 4-3 (left), we obtain from the slope of the bias term (again using a least-squares fit) for $\tau \leq 1.5$, the rate $\gamma_1 \sim 1.3$. Therefore, we obtain $r = 1.3/0.85 = 1.5$ as the rough estimate needed to determine the rate of convergence under our model

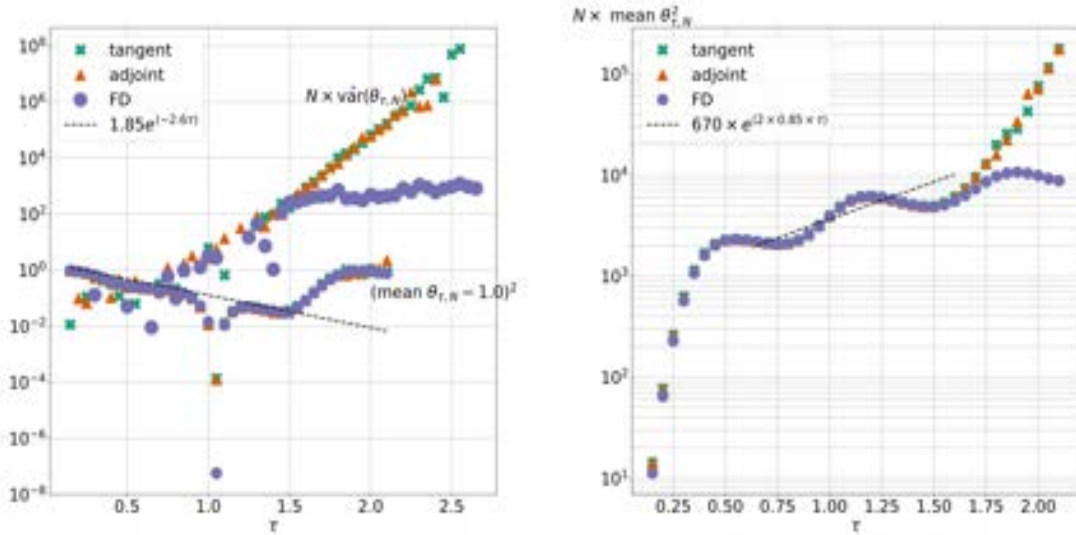


Figure 4-3: Left: Estimates of the variance and the bias of $\theta_{\tau,N}$ as a function of τ , for the Lorenz '63 system outlined in section 4.4.1. The dashed line indicates the least-squares fit over $\tau \leq 1.5$. Right: Sample mean estimates of $\mu_s(\theta_{\tau,N}^2)$ as a function of τ for the Lorenz '63 system outlined in section 4.4.1. The dashed line indicates the least-squares fit over $\tau \leq 1.5$.

assumptions in sec 4.3.3. Thus, we note from 4-2 that $\beta \sim 0.5$. It would not be gainful to also estimate k since we only seek an upper bound for β which is quite insensitive to k , in the first place. We can interpret this rate as the best possible rate of convergence for the Lorenz '63 system.

Agreement with Eyink *et al.*'s results

Eyink *et al.* [96] propose that the probability distribution of $\theta_{\tau,N}$ for the Lorenz '63 system is a fat-tailed distribution and does not obey the classical CLT. Consequently, our assumption of bounded variance of $\theta_{\tau,N}$ for finite values of τ should fail for this system. Our numerical results indicate that both the variance and bias trends are worse than our optimistic model, thus confirming their predictions. However, locally around $\tau \sim 1$, operating under the assumption that the bias and the variance trends can be empirically modelled using our optimistic estimates, we were able to provide an upper bound on the actual convergence rate. Owing to the fact that $\tau \sim 1$ is not large enough for the failure of the CLT assumption to be manifest, our empirically

determined rate of 0.5 provides an upper bound on that estimated using Eyink *et al.*'s tail estimates. Thus we conclude firstly that our numerical results provide further evidence in support of the failure of the CLT for $\theta_{\tau,N}$. Secondly, empirical determination of the optimistic rate of convergence under the CLT assumption locally around $\tau \sim 1$, predicts an upper bound at the rate of a Monte Carlo simulation, confirming Eyink *et al.*'s observations.

It is worth commenting on the fundamental difference between the distribution of sample averages of state functions and that of $\theta_{\tau,N}$, which is a sample average of the derivative of a state function (averaged over a short time). The former distribution has bounded variance but the latter is not guaranteed to, as this example demonstrates. In fact, it has been shown [136] that N -sample averages (computed numerically as ergodic averages) of state functions converge at the rate of $1/\sqrt{N}$. However, the behavior of the N -sample average that is $\theta_{\tau,N}$, is unlike that of state functions. $\theta_{\tau,N}$ converges to $\mu_s(\theta_{\tau,N})$ (due to the law of large numbers) at a rate slower than $1/\sqrt{N}$. This rate decreases as τ increases, bearing on the poor accuracy of the estimate $\theta_{\tau,5 \times 10^6}$ for values of $\tau \gtrsim 1.5$ seen in 4-3 (left), despite using a seemingly large number of samples.

Another observation that can be made from 4-3 (left) is that unlike the variances of $\theta_{\tau,N}^A$ and $\theta_{\tau,N}^T$, $\text{var}(\theta_{\tau,N}^{\text{FD}})$ appears to saturate for $\tau \gtrsim 1.5$. This is explained by the fact that $|\theta_{\tau,N}^{\text{FD}}|$ is bounded above by $\|J\|_\infty/\epsilon = c/\epsilon$, where ϵ is the value of the parameter perturbation in the finite difference approximation. The value c is the supremum of the z coordinate of the Lorenz attractor, which is a finite value since the attractor is a bounded set. The fact that $\theta_{\tau,N}^{\text{FD}}$ is bounded for all τ can also be observed in the saturation of the estimate of $\mu_s(\theta_{\tau,N}^{\text{FD}^2})$ shown in 4-3.

4.4.2 The Lorenz '96 model

Our results for the Lorenz '63 system in section 4.4.1 show that an upper bound on the rate of convergence for an optimal choice of τ (~ 1) is 0.5. Although the actual rate is slower than a typical Monte Carlo simulation, one can immediately see that since the system is low-dimensional, ES is still practical in the Lorenz '63 system.




Figure 4-4: The variance and the square of the bias in the ensemble tangent sensitivity estimator for the Lorenz '96 system. The dashed line indicates the least-squares fit of the variance data.

In this section, we aim to assess the computational practicality of the ES method in a higher-dimensional system. We consider the atmospheric convection model, called the Lorenz '96 [180] model, that is known to have a strange attractor [152]. The 40-dimensional state vector $u := [x_1, \dots, x_{40}]^T$ evolves according to the following set of odes [180, 152]

$$\frac{dx_k}{dt} = (x_{k+1} - x_{k-2})x_{k-1} - x_k + s, \quad k = 1, \dots, 40, \quad (4.25)$$

where, s , the parameter of interest, denotes an external driving force. The first term on the right hand side of equation 4.25 represents nonlinear advection and the second term represents a viscous damping force. The components of the state vector are periodic in the sense that $x_k = x_{40+k}$, $k \in \mathbb{N}$. We define our objective function $J(u(t))$ to be the mean of the components of u , i.e.,

$$J(u(t)) = (1/40) \sum_{k=1}^{40} x_k(t).$$

We use the value $s = 8.0$ for the forcing term, at which the system has been shown to exhibit chaotic behavior [152]. Time integration of the primal system of equations 4.25 is performed using a fourth order Runge-Kutta scheme with a timestep of 0.01 time units. The distribution of each component of the state vector for this system is

known from the literature [267] to converge to a Gaussian and the time average of J approximately becomes a linearly varying function with s , on long-time evolution. The slope of $\mu_s(J)$ vs. s estimated from our computations and previous work [152] is 0.096.

In Fig. 4-4, we report the squared bias and the variance of the estimator $\theta_{\tau,N}^T$ as a function of τ . The total cost, T , is set at 10^4 time units across different values of τ . The values of $\mu_s(\theta_{\tau,N})$ and $\mu_s(\theta_{\tau,N}^2)$ are computed as sample means and used in the approximation of the bias and variance (denoted by $\hat{\text{var}}(\theta_{\tau,N})$ in Fig. 4-4 to indicate the approximation as a sample mean) terms, identical to our description in section 4.4.1 for the Lorenz '63 system. From Fig. 4-4, we can see that the variance follows our model assumptions in section 4.3.2; we find that the variance is exponential in τ at the rate ≈ 4 , which is close to twice the leading Lyapunov exponent reported in the literature [152, 267]. From Fig. 4-4, ignoring initial transients, it can be seen that the bias term does not show an exponential convergence, not even locally in the vicinity of $1/\lambda_1 \approx 0.5$. As a result, our local analysis to produce a rough estimate under our model assumptions, as we did in section 4.4.1, is not applicable in this case. This implies that the rate of convergence is worse than our model in 4.3.2 (for any r) since the bias falls slower than the assumed exponential. One can argue that the mean squared error (sum of the squared bias and the variance) is low in absolute value for τ near 0.5 deeming the ES estimator to be reasonably accurate and therefore practically applicable, although the asymptotic rate of convergence in τ may be low. However, note that the rate of convergence is an objective function-independent measure of the ES estimator – the mean squared error may be coincidentally within required accuracy for a choice of τ , for this particular objective function. Thus, it is reasonable to conclude that the ES method is infeasible for the Lorenz '96 model.

4.4.3 Chaotic flow over an airfoil

In this example, we discuss the numerical simulation of an unsteady, chaotic flow around a two-dimensional airfoil. So far, we have assessed the convergence of the ES method by observing the trends in the bias and variance of $\theta_{\tau,N}$ in low-dimensional




Figure 4-5: Tangent density field v_ρ at time $t = 78 c/a_\infty$.

systems. Our numerical results were informative enough to predict the rate of convergence while simultaneously being within the limits of practical computation, owing to the low dimensionality of the systems considered in sections 4.4.1 and 4.4.2. In contrast, in a typical chaotic CFD simulation, it would not be practical to numerically estimate the bias and variance trends of $\theta_{\tau,N}$. We will thus attempt to predict the convergence trend using a single finite-difference solution (that can be used to compute one sample of $\theta_{\tau,N}$). Our goal is to use a physics-based approach that eliminates the need for a rich $\theta_{\tau,N}$ dataset. We consider the NACA 0012 airfoil at the Reynolds number $Re_\infty = 2400$ and Mach number $M_\infty = 0.2$ at an angle of attack $\alpha = 20^\circ$. Although the flow physics in three-dimensional turbulent flows is more complex, the two-dimensional airfoil case we consider exhibits the phenomena of stall and flow separation that are responsible for the chaotic behavior. For an extensive analysis of the Lyapunov spectrum and its dependence on the numerical discretization for this problem, see [101, 222]. The primal system is the set of compressible Navier-Stokes equations, which is discretized in space using a third-order Hybridizable Discontinuous Galerkin (HDG) method [100, 99] with the Lax-Friedrichs-type




Figure 4-6: Time evolution of L^2 norm of the tangent fields corresponding to the density ρ (top) and x -momentum ρu (bottom).

stabilization matrix in [98]. The computational domain spans $\sim 10 - 20$ chord lengths away from the airfoil and is partitioned using an O-grid with 35280 isoparametric triangular elements. The use of a large computational domain is customary in external aerodynamics to reduce the change in the effective angle of attack induced by the missing vortex upwash. The state vector u is defined to be the set of coefficients in the basis function representation of the density, the components of the momenta and the energy, at all elements in the computational domain. Our objective function J is the lift coefficient and the parameter of interest s is the freestream Mach number, M_∞ . A three-stage, third-order diagonally-implicit Runge-Kutta (DIRK) method [5] is used for the temporal discretization. A no-slip, adiabatic wall boundary condition is imposed on the airfoil surface, and the characteristics-based, non-reflecting boundary condition in [100, 99] is used on the outer boundary. The spin-up time before the sensitivity computation is performed is $t = 10,000 c/a_\infty$, where a_∞ denotes the far-field speed of sound and c is the chord length. We note that the chosen spin-up time is one order of magnitude larger than the time required for the convergence of time-averaged flow quantities [101]. At the spin-up time, we reset $t = 0$ and all the

results below are indicated with respect to this new initial time.

Since a tangent solver is not available typically in CFD codes, the tangent perturbation fields denoted by $v(t)$ are computed non-intrusively with regard to the primal solver using a finite difference approximation (FDA). The parameter perturbation has a magnitude of $\epsilon = 10^{-6}$. Fig. 4-5 shows the sensitivity field of the fluid density with respect to the perturbation in the freestream Mach number at time $t = 78c/a_\infty$. This quantity, denoted in Fig. 4-6 as v_ρ , can be used to compute a single sample of $\theta_{\tau,N}^{\text{FD}}$. The airfoil, much smaller in size than the computational domain, is annotated in Fig. 4-5. Entropy and acoustic waves are generated on the outer boundary at $t = 0$ due to the M_∞ perturbation. From the entropy wavefront shown in Fig. 4-5, one observes that the entropy wave does not reach the airfoil yet at $t = 78c/a_\infty$. The time evolution of the L^2 norm of the tangent fields corresponding to the density, v_ρ , and the x -momentum, $v_{\rho u}$, are shown in 4-6. As we discussed in section 4.4.1, finite difference sensitivities tend to saturate with time as opposed to tangent and adjoint sensitivities which keep increasing exponentially. This is due to the finite difference sensitivities having an upper bound proportional to $1/\epsilon$, since the supremum of the objective function over phase space is a finite value independent of time. As indicated in Fig. 4-6, the entropy wave reaches the airfoil at $t \sim 115c/a_\infty$. By this time, the sensitivity fields already saturate. We can estimate by extrapolation that $\|v_\rho\|, \|v_{\rho u}\| > 10^5$ if computed using the tangent equation, at $t = 115c/a_\infty$.

One expects the bias in the ES estimator to be non-negligible for trajectory lengths shorter than $t = 115c/a_\infty$. That is, the sensitivities must be computed at least until the time the information about the perturbation propagates to the airfoil, in order for the average sensitivity of the lift to converge to its true value. In other words, convergence of the bias requires the entropy wave to reach the airfoil and thus τ must be larger than $\tau^* := 115c/a_\infty$. In order to predict the cost of the ES method, let us make the optimistic assumptions that at τ^* , the bias term is close to zero and that the CLT holds for the variance. A random tangent field is at least $\sim 10^5$ in magnitude at τ^* implying that the variance at τ^* would be $\mathcal{O}(10^{10})$. Therefore, under the CLT assumption, which we have shown to be too optimistic in our previous




Figure 4-7: Left: Adjoint field corresponding to the density, labelled rhoa in the colormap, at $t^* = 0.35$. Right: The L^2 norm of the adjoint vector field, y , for the Navier-Stokes system in section 4.4.4 as a function of time.

examples in sections 4.4.1 and 4.4.2, one would require on the order of 10^{10} samples for an $\mathcal{O}(1)$ mean squared error. Solving the primal and the perturbation equations on the order of 10 billion times for trajectories of lengths τ^* would be computationally infeasible. This example illustrates that the much shorter timescale of divergence of the adjoint/tangent fields, compared to that of the convergence of the bias, leads to computational intractability of the ES method.

4.4.4 Turbulent flow over a turbine vane

As a final example, we consider an implicit LES of turbulent flow around a highly loaded turbine nozzle guide vane performed by Talnikar *et al.* [257]. Our simulations approximate the aero-thermal experimental investigations of turbine guide vanes in a linear cascade arrangement at the von Karman Institute for Fluid Dynamics [17]. The linear cascade is approximated in the simulation domain using periodic boundary conditions in the transverse and spanwise directions. The spanwise extent, is restricted to 0.15 times the chord length, which has been numerically verified to be sufficient to capture the turbulent flow physics. The Reynolds number of the flow is 10^6 . The isentropic Mach number computed with respect to the static pressure at the outlet is 0.9. Isothermal boundary condition is used on the surface of the vane. The

fluid achieves transonic speeds as it flows over the surface of the vane, transitioning to turbulence on the suction side.

The primal problem is a compressible Navier-Stokes system as in 4.4.3, solved here using a second order finite volume scheme [257]. A strong stability-preserving third order Runge-Kutta scheme is used for time integration and a weighted essentially non-oscillatory scheme [173] is used for shock capturing. The mesh is generated by uniformly extruding a two-dimensional hybrid structured/unstructured mesh in the spanwise direction.

The long-time averaged objective of interest is the mass-flow averaged pressure loss coefficient on a plane 0.25 chord lengths downstream of the trailing edge of the vane. The reduction in stagnation pressure loss is due to mixing in the turbulent wake and the formation of a boundary layer on the surface of the vane. Sensitivities (computed using the adjoint method) are calculated with respect to Gaussian-shaped source term perturbations to the system dynamics centered 0.3 chord lengths upstream from the leading edge of the vane in the axial direction. The time variable, t^* , is nondimensionalized with respect to the flow-through time, which is the duration of the time taken for fluid flow from the inlet to outlet. In 4-7 (left), we show the adjoint field corresponding to the density, denoted by ρ_a , at $t^* = 0.35$. 4-7 (right) shows the L^2 norm of the adjoint field, denoted by y , as a function of t^* . From the growth of the adjoint vectors in 4-7, we can estimate the leading Lyapunov exponent for the flow to be ~ 46 in nondimensional units. The time taken for the entropy wave of the adjoint solution to reach the source of the perturbation (the propagation direction of the adjoint solution is reversed with respect to the primal flow) is $\gtrsim 0.5$ time units. From 4-7, we can see that the adjoint solution has diverged by 10 orders of magnitude by this time.

The reason for the rapid divergence is the high angle of attack of the flow onto the vane surface which causes the Mach wave of the adjoint solution to propagate with a maximum speed of Mach 1.9 upstream in the axial direction. The entropy wave, in comparison, has a lower maximum speed of Mach 0.9, which is equal to the flow speed. The adjoint diverges quickly once the Mach wave reaches the turbulent boundary layer

region close to the trailing edge of the vane on the suction side, as shown in 4-7. But the bias in the adjoint computed over an ensemble of trajectories can be expected to decline only for trajectory lengths $\gtrsim 0.5$ time units, which is approximately the time required for the entropy wave to reach the vane leading edge. Similar to the flow over the airfoil presented in section 4.4.3, the lower speed of propagation of the entropy wave from the inlet makes the ES method infeasible in this problem.

Equipped with the estimate of the Lyapunov exponent and the timescale of convergence of the bias, we predict the cost of the ES method in this case. As in section 4.4.3, let us assume the best case scenario for both the bias and the variance terms. Suppose the bias has negligible magnitude at 0.5 time units and that the CLT assumption is valid. Since the true variance would be on the order of 10^{20} at $t^* = 0.5$, we would require $\mathcal{O}(10^{20})$ samples in order to reduce the variance of $\theta_{0.5,N}^A$ and consequently, the mean squared error to $\mathcal{O}(1)$. To roughly estimate the computational power required, we need about 10 teraFLOPS for 12 hours for a single adjoint simulation up to $t^* = 0.5$; this means approximately 10^{15} exaFLOPS for 12 hours would be needed for convergence. This is beyond the computational capabilities achievable in the near future. To conclude our discussion, even though the source of perturbations is close to the leading edge of the vane, turbulence at a high Reynolds number makes the timescale of the growth of perturbations much shorter than the time required for the propagation of the information about the perturbation through entropy waves. As a result, since the number of samples required for convergence increases exponentially at best with time, the ES method becomes computationally intractable.

Before closing, we remark on the differences between this example and the previous one in section 4.4.3. Although the timescale discrepancy argument supplied to rule out the applicability is the same in the two examples, the appearance of such a discrepancy is attributed to rather different reasons. In the airfoil flow, one could have argued that the particular choice of the parameter perturbation prolonged the timescale for the decay of the bias, leading to the divergence of the tangent/adjoint within that timescale. That situation is typical in an external flow, where parameters we can influence are at the farfield, spatially separated from quantities of interest.

The flow over the turbine vane in this section, on the other hand, is an internal flow where the flow physics (the slowness of the mean flow compared to the Mach wave) was responsible for the inapplicability of the ES method, rather than the limited choice of controllable parameters.

4.5 Discussion and comments

From the optimistic analysis presented in section 4.3 and the numerical examples considered in the previous section, one can conclude that the ES method cannot be a universal chaotic sensitivity analysis method, owing to its computational cost. The main basis for this conclusion is that the ES estimator's bias often decays at a longer timescale than the inverse of the largest Lyapunov exponent, which dictates its variance. The analyses and the examples do not preclude however, cases where this timescale discrepancy is not large enough to make the ES method computationally impractical. This leads to the question of whether one can choose specific objective functions and parameters, that are dissimilar to the examples in section 4.4, so as to yield better scenarios for the application of the ES method. Here we provide a closer examination of this difficult and problem-specific question; we draw upon connections between signature timescales, that could guide us in the determination of the edge cases for the method's applicability.

As noted in section 4.3, in uniformly hyperbolic systems that we consider here, the decay of correlations is exponential. Using the Collet-Eckmann conjecture, we argued that the rate of the bias decay ($= \gamma_1$) is similar to the rate of correlation decay. This argument is also exemplified, in addition to our heuristic explanation considering only stable perturbations in section 4.3.2, through alternative derivations of Ruelle's response formula that have appeared in a few works in the dynamical systems literature ([49, 28]; see also Chapter 7 or proposition 8.1 of [118] for a computable modification of Ruelle's formula in discrete-time systems). The thrust of these alternative formulations is that one can replace the computation of Ruelle's formula which, in its original form (Eq. 4.8) is a time integral with an exponentially increasing variance,

with a computation that has bounded variance at all times. These formulae have the convergence of a typical Monte Carlo estimator (at the rate of $1/\sqrt{N}$), independent of the objective function, and the bias in these formulae decay as the time correlations in the system. This implies that, supposing $\theta_{\tau,N}$ behaved like a Monte Carlo estimator for a specific objective function (or only slightly worse, as in the Lorenz '63 case), the bias decay rate is still bounded above by the correlation decay rate.

It seems unintuitive that there exists a connection between the two timescales. After all, the two exponential rates, of the decay of correlations and that of the Lyapunov exponents, represent disconnected chaotic phenomena, namely, the rate of loss of information carried by state functions, and local expansion of infinitesimal perturbations, respectively. Nevertheless, several works have verified this connection in low-dimensional uniformly hyperbolic dynamics (see [246, 198] for recent work). In fact, numerical studies in disparate turbulent flows such as drift wave turbulence in plasma [217] and in fluid convection [244], also confirm that the correlations decay slower than $(1/\lambda_1)$. Given the validity of the chaotic hypothesis in many turbulent flows, there is reason to expect this connection between the two timescales, that can be rigorously shown to hold only in mathematical toy models, to also manifest itself in general turbulent flows.

At the same time, the rate of correlation decay depends on the observables considered; this is why one cannot exclude the existence of objective function-parameter pairs that lead to fast convergence of $\theta_{\tau,N}$, even though the decay of correlations for the system as a whole is at a slower rate than λ_1 (i.e., even if the above conjecture holds). The decay rate of correlations for a chaotic system is generally determined as the average decay between two functions belonging to a broad function class (such as the set of continuously differentiable functions of the state). There exist observables correlations between which decay at a faster rate than that given by this signature rate of the system. The correlation time at which the bias converges is rather difficult to predict a priori. In accordance with the approach of this work which isolates extreme scenarios, we can take the autocorrelation times of the worst-case observables (or the ones with notably slow decay rates) as a crude approximation of the upper

bound on the bias decay time. In the event that this upper bound is smaller than the inverse of the largest Lyapunov exponent, we can expect the ES method to behave at least as well as a Monte Carlo estimator.

In aeroacoustics, an important application area where RANS modeling is inadequate and turbulence-resolving LES/DNS is resorted to, farfield noise (away from a nozzle) caused by a turbulent jet is often used as an objective function. In such a scenario, a wealth of spectral studies (for brevity, we refer only to a few recent works on modelling turbulent jets [241, 43, 105]) have found that the energy-carrying spatial Fourier modes (of lower wave number) appear as large-scale structures with high spatial coherence. This also translates into high coherence in time or equivalently, a slow decay of time correlations of streamwise observables – this being corroborated by more sophisticated mechanisms of turbulent structure transit than the oft-used Taylor’s hypothesis of frozen turbulence (among the vast body of work on space-time correlations in turbulent flow, we point the reader to some comprehensive reviews in [132, 230, 137] that detail this point). So, in this case, we can take the integral timescales reported for farfield streamwise measurements as the worst-case estimate for the bias decay time. Similarly, in turbulent boundary layers, integral timescales associated with the intermittent or outer regions of the flow can be taken as an upper bound since very near-wall regions exhibit faster decay of correlations. In the case that the slowest rate of correlation decay estimated as the inverse of this upper bound, is still larger than the largest Lyapunov exponent, the ES method could be practical.

Given an objective function of interest and a fixed set of parameters, one might in practice, benefit from knowing the decay of correlations from thoroughly documented computational DNS/LES datasets or experimental observations for the appropriate type of flow (for example, see [105, 241, 43] for the case of jets and [202, 36, 20] for the case of turbulent boundary layers). In this regard, the availability of space correlation data, which are more naturally obtained and commonly reported in computational studies, should also prove to be sufficient for estimation purposes since they can be transformed approximately to time-correlation data. If an estimate of the largest Lyapunov exponent is also available, one needs to simply compare the timescales and

use the results of section 4.3 to predetermine if the ES method would be applicable.

Ultimately, these heuristic predictions based on prior knowledge of timescales can only help us identify the corner cases of applicability (or the lack thereof) of the ES method. We contend that it is more likely to encounter undecidable cases and close this discussion with such an example that appears in Rayleigh-Bénard convection. Sirovich *et al.* [244] have shown that the correlation time defined as the time it takes for the streamwise velocity autocorrelations to drop to its first minimum, is on the same order of magnitude as the largest Lyapunov exponent. Interestingly, our numerical results for the Lorenz '63 system (section 4.4.1), which is a mathematical model for Rayleigh-Bénard convection, is also in line with this finding – the rate of decay of the bias and the largest (and only) positive Lyapunov exponent are both $\mathcal{O}(1)$. In such a scenario, the practitioner may choose to rely on alternative methods (such as [206, 185, 37] or Chapter 7) in order to avoid incurring the cost of solving an enormous number of tangent or adjoint equations without being to establish convergence.

4.6 Conclusion

To compute sensitivities with respect to design parameters, of statistically stationary quantities in chaotic systems, ensemble methods appear to be an appealing solution; they are both conceptually simpler and easier to implement than fluctuation-dissipation-based [185] and shadowing-based [211] methods, all of which are, moreover, still under active development. However, the present work has shown that Eyink *et al.*'s [96] results revealing poor convergence of ensemble computations in the Lorenz '63 system, is more widely representative of convergence trends in general chaotic systems. The present work deals with estimating a theoretical upper bound on the rate of convergence of an ensemble sensitivity, agnostic to the objective function-parameter pair.

For this, the most optimistic assumptions are made on the bias and variance associated with the ensemble sensitivity (ES) estimator, under the mathematical simplification of uniform hyperbolicity. We show that, with the integration time, in the

best case, the bias decays exponentially at a problem-dependent rate and the variance increases exponentially at the rate of twice the largest Lyapunov exponent. Assuming these optimistic bounds, the computational cost of the ES method, in theory, still scales exponentially with the mean squared error. The ratio of the bias decay rate to the leading Lyapunov exponent of the system, is the single most influential parameter that determines the feasibility of the method.

Our numerical results for the Lorenz '63 system show that the optimistic model proposed for the least mean squared error is only locally applicable. The upper bound on the rate of convergence is 0.5, concurring with Eyink *et al.*'s [96] results. The 40-variable Lorenz '96 system serves as an example of a low-dimensional attractor for which the asymptotic convergence is remarkably slow. Although the rate of convergence is poor for this system, the mean squared error magnitudes were low at a reasonable computational cost, for the chosen objective function. This suggests that one may encounter, in practice, objective functions for which the ensemble sensitivities are within a specified accuracy, at an affordable computational cost. In the numerical simulations of chaotic fluid flow we consider, we obtain optimistic estimates on the rate of convergence which hold true for a general objective function. Our results indicate that the flow physics imposes an upper bound on the rate of convergence. Altogether, the present numerical evidence suggests the following: even under the optimistic assumption of exponential decay of the bias, the cost of exponential sampling of an expensive primal problem can make the ES method infeasible in practical applications, for a general objective function-parameter pair. However, there may exist objective function and parameter choices that lead to a smaller timescale discrepancy between the bias decay and the variance increase – this could lead to a faster convergence than the estimate of the upper bound.

Chapter 5

Computing linear response to unstable perturbations

In this chapter, we present a computable reformulation of Ruelle’s linear response formula for chaotic systems. The new formula achieves an error convergence of the order $\mathcal{O}(1/\sqrt{N})$ using N phase points. The reformulation is based on splitting the overall sensitivity into that to stable and unstable components of the perturbation. The unstable contribution to the sensitivity is regularized using ergodic properties and the hyperbolic structure of the dynamics. Numerical examples of uniformly hyperbolic attractors are used to validate the formula against a naïve finite-difference calculation; sensitivities match closely, with far fewer sample points required by the new formulation.

This chapter contains an early version of the space-split sensitivity or S3 algorithm, to be introduced in the next chapter. This pre-S3 version computes linear response exactly when a) the chaotic system is purely expanding and there is no stable manifold and b) in rare [130](see also Chapter 19 of [154]) uniformly hyperbolic attractors in which stable and unstable subspaces are differentiable; typically these subspaces are only Hölder continuous on the phase space. The pre-S3 algorithm presented in this chapter yields accurate results in uniformly hyperbolic systems with one-dimensional unstable manifolds that obey these restrictive assumptions. Moreover, the treatment of the unstable contribution in this chapter is an alternative to the same in Chapter

7, which contains a more general version of the S3 algorithm. Since the assumption of differentiability leads to an easier formulation, we treat systems with arbitrary-dimensional unstable manifolds in this chapter, in contrast to the next two chapters, which deal with one-dimensional unstable manifolds. This chapter is also available online as a preprint at [62]. Finally, if reading the thesis in order, the reader may wish to keep the introduction, which has been retained to keep this chapter self-contained.

5.1 Introduction

Given a dynamical model, how much do the outputs change in response to small changes in input parameters? The computational aspect of this common question, which arises across science and engineering, entails the mature discipline of sensitivity analysis. Using numerical simulations or experimental observations of dynamical systems, the computed responses or sensitivities have enabled multidisciplinary design optimization, uncertainty quantification and parameter estimation in diverse non-chaotic models (see [219, 52] for recent surveys in aerodynamic systems). Sensitivity analysis in chaotic systems, however, still remains nascent. This is because traditional sensitivity computation is done through linear perturbation methods including tangent or adjoint equations, and automatic differentiation, but these are inherently unstable in chaotic systems [206, 271]. More sophisticated techniques are needed to calculate the long-term effects of parameter changes in chaotic systems. Some of them are being actively investigated as of this writing, and face challenges such as lack of convergence guarantees or prohibitive computational cost. In this work, we develop an alternative method¹ for addressing some of these challenges. We propose a formula for differentiating statistical averages in chaotic systems to parameter perturbations, in a way that is amenable to computation. We postpone until section 5.2 the precise definition of the statistical response we aim to compute and the desired qualities of a computational solution.

Currently, the more sophisticated approaches to chaotic sensitivity computation

¹a preliminary version of the derivation of the method has appeared in [60]

include a) the ensemble sensitivity method b) shadowing-based approaches and c) perturbation methods on the transfer operator. The non-intrusive least squares shadowing (NILSS) method (see [37, 211], and [271] for the different versions of NILSS) computes a shadowing perturbation that remains bounded in a long time window under the tangent dynamics. However, the sensitivity computed by using the shadowing tangent solution is not guaranteed to be an unbiased estimate of the true sensitivity [38]. This is because while ergodic averages converge for almost every trajectory, there are measure zero subsets of the attractor (e.g. unstable periodic orbits [120]) on which they do not converge. We therefore seek an alternative that does not rely on computations along a single trajectory that is not guaranteed to be typical. The Lea-Allen-Haine ensemble sensitivity method [96] suggests a work-around to the exponentially diverging sensitivities computed by the conventional tangent/adjoint methods, by truncating the tangent/adjoint equations at a short time, and taking a sample of average of many such short-time sensitivities. But, although these sample averages converge in the infinite time limit, they are prohibitively expensive because the variance of tangent/adjoint solutions increases exponentially with time ([69]).

In a recent work, Crimmins and Froyland [81] have developed a new Fourier analytic method for constructing the Sinai-Ruelle-Bowen or SRB measures, the invariant probability distributions perturbations of which we are interested in, of uniformly hyperbolic dynamics on tori. They construct a matrix representation of perturbed transfer operators that are quasi-compact on certain anisotropic Banach spaces (see [118] for uniformly hyperbolic systems in particular, and [28] for a recent review). Using this matrix representation, the leading eigenvector, which is the SRB measure, is then approximated. Although equipped with a strong theoretical basis, this method would be overkill for our purpose, since we do not need to construct the SRB measure, but only compute the sensitivity of a given expectation with respect to it. Moreover, methods based on perturbations of the transfer operator such as [81, 185] typically have a computational cost that scales poorly with the problem dimension since they either involve Markov partitions (specific discretizations of the attractor) or need a number of basis functions to approximate the eigendistribution (the SRB measure)

that scales with the dimension. Another recent method [166] computes sensitivities by solving an adjoint equation as a boundary value problem on periodic orbits, although questions surrounding the convergence of the computed sensitivity to the true sensitivity (like in the shadowing methods) must be investigated further.

The strategy developed in this chapter deviates from that of all the above-mentioned methods. However, like ensemble sensitivity methods, it builds upon Ruelle’s formula. While ensemble sensitivity suffers from the unbounded variance of the unstable contribution to the overall sensitivity, our new formulation splits the contributions and performs a finite-sample averaging of tangent equation solutions only for the stable contribution. The unstable contribution is manipulated through integration-by-parts and using the statistical stationarity (measure preservation) of the system to yield a computation that does not use unstable tangent solutions. Since both parts of the sensitivity are computed through sampling on generic flow trajectories, the problem of the computed sensitivities corresponding to atypical trajectories, which shadowing-based methods are vulnerable to, is averted. The chapter is organized as follows. In the next section, we give the problem setting and state results from dynamical systems theory that are used in the derivation of pre-S3. Our main results are stated in 5.3, and the derivation of the pre-S3 formula follows in 5.4. An interpretation of the unstable contribution, as derived in pre-S3, is presented in 5.5, and some comments on its computational details can be found in 5.6. Numerical examples demonstrating a naïve implementation of the pre-S3 formula are reported in 5.7, and the conclusions follow in 5.8.

5.2 Problem statement

In this section, we define the output quantity of interest for pre-S3 computation. Where they are used, we provide a brief description of concepts from ergodic theory and dynamical systems, in order to make this a self-contained presentation for computational scientists from different fields.

5.2.1 The primal dynamics

Our primal system is a C^3 diffeomorphism φ^s , on a d -dimensional compact manifold M , parameterized by a set of parameters s ,

$$u_{n+1} = \varphi^s(u_n), \quad n \in \mathbb{Z}, u_n \in M. \quad (5.1)$$

In a numerical simulation, $u_n \in M$ represents a solution state (a d -dimensional vector) at time n and the transformation φ^s can be thought of as advancing by one timestep. For simplicity, s is assumed to be a scalar. Without loss of generality, we take $s = 0$ to be the reference value and the map φ^0 is simply written as φ , without the superscript. The state vector $u_n, n \in \mathbb{Z}^+$ is a function of the initial state u_0 ; explicitly, $u_n = \varphi_n(u_0)$, where the subscript $n \in \mathbb{Z}^+$ in φ_n^s refers to an n -time composition of φ^s , and φ_0^s is the identity map on M . We also use the notation $\varphi_{-n}, n \in \mathbb{Z}^+$ to denote the inverse transformation $(\varphi)^{-1}$ composed with itself n times; that is, $\varphi_{-n}(u_n) = u_0$.

5.2.2 Ensemble and ergodic averages

We assume that the dynamics φ^s preserves an ergodic, physical probability measure μ^s of an SRB-type (see [281] for an introduction to SRB measures), which gives us a statistical description of the dynamics. In particular, expectations with respect to μ^s can be observed as infinitely long time averages along trajectories: if $f \in L^1(\mu^s)$ is a scalar function, then $\lim_{N \rightarrow \infty} (1/N) \sum_{n=0}^{N-1} f(\varphi_n^s(u)) = \langle f, \mu^s \rangle, u$ Lebesgue-a.e on the basin of attraction. The superscript s in μ^s emphasizes the dependence of the SRB measure on the parameter; simply μ , without the superscript, refers to the SRB measure at $s = 0$. The ensemble average or the expectation of f with respect to μ^s , $\langle f, \mu^s \rangle$, is a distributional pairing of f with μ^s : the integral of f on the phase space weighted by μ^s . We sometimes use the shorter notation $\langle f \rangle^s$, and without the superscript at $s = 0$, to denote the same quantity. The infinite time average, called the ergodic average, is the more natural form of $\langle f \rangle^s$ from the computational/experimental standpoint, since it can be obtained by numerical evaluation/measurements of f along trajectories. In practice the ergodic average is computed up to a large N and this is used to

approximate the ensemble average $\langle f \rangle^s$ – by ergodic average, we mean this long but finite time average, in the remainder of this work.

5.2.3 Quantity of interest

We are interested in determining the sensitivity of the ensemble average of an objective function $J \in \mathcal{C}^2(M)$, $\langle J \rangle^s$, to s : $d_s \langle J \rangle^s = \langle J, \partial_s \mu^s \rangle$. The regularity of J , along with other assumptions on the dynamics which we detail where they appear, are such that the linear response formula of Ruelle [232, 234] is satisfied. The objective function J can also explicitly depend on s , and assuming that J is continuously differentiable with respect to s , the quantity of interest $d_s \langle J \rangle^s = \langle \partial_s J, \mu^s \rangle + \langle J, \partial_s \mu^s \rangle$. As we will see however, in a chaotic system, the mathematical or algorithmic difficulty lies in computing the derivative of the SRB measure in the second term, and not in the first, which can be computed as an ergodic average, assuming the function $\partial_s J(u, s)$ is known. Thus from now on we ignore the first term and develop an algorithm for $\langle J, \partial_s \mu^s \rangle$, which is nontrivial to compute – we describe precisely why, in the next section.

5.2.4 Ruelle’s formula and its computational inefficiency

To introduce Ruelle’s formula, let the matrix-valued function $D\varphi : M \rightarrow GL(d)^2$ give us the Jacobian matrix of the transformation φ . Here D refers to the derivative with respect to phase space; so, the Jacobian at u , $D\varphi(u)$, is a map from the tangent space of M at u , denoted as $T_u M$, to $T_{\varphi(u)} M$. We now introduce a more succinct notation for the tangent operator in the following definition.

Definition 1. *The tangent operator $\mathcal{T}(u, n) : T_u M \rightarrow T_{\varphi_n(u)} M$ is a linear operator (a matrix) for each $n \in \mathbb{Z}$ and is defined as the derivative of φ_n with respect to the state vector u , evaluated at u . By this definition, it can be written as Jacobian matrix*

² $GL(d)$ is the set of all invertible matrices of dimension $d \times d$.

products, as follows,

$$\mathcal{T}(u, n) := \begin{cases} D\varphi(\varphi_{n-1}(u)) \cdots D\varphi(u), & n > 0 \\ (D\varphi(\varphi_n(u)))^{-1} \cdots (D\varphi(\varphi_{-1}(u)))^{-1}, & n < 0 \\ \text{Id}, & n = 0. \end{cases}$$

From the linear response theory that was rigorously developed by Ruelle [232, 234], we have the following formula for the sensitivity of interest,

$$\left. \frac{d\langle J \rangle^s}{ds} \right|_{s=0} = \sum_{n=0}^{\infty} \langle D(J \circ \varphi_n(\cdot)) \cdot X(\cdot), \mu \rangle, \quad (5.2)$$

where $X(u) := \partial_s \varphi^s(\varphi^{-1}(u))$ denotes the vector field corresponding to parameter perturbation. Here DJ refers to the derivative of J with respect to the phase space. For brevity, we adopt the subscript notation also with scalar and vector fields, as explained below:

1. if $f : M \rightarrow \mathbb{R}$ is a scalar field, f_n is used to denote the function $f \circ \varphi_n$.
2. The derivative of $f_n = f \circ \varphi_n$, evaluated at a μ -typical point u , is denoted as $Df_n(u) := D(f \circ \varphi_n)(u)$. On the other hand, $(Df)_n$ refers to the derivative of f evaluated at $u_n = \varphi_n(u)$, where again u is a μ -typical point.
3. if V is a vector field, then $V_n(u) \in T_{\varphi_n(u)}M$ denotes its value at u_n : $V_n(u) := V(u_n)$.

Using this notation, the integrand in the n th summand of Ruelle's formula (Eq. 5.2) can be written as $D(J \circ \varphi_n) \cdot X = DJ_n \cdot X$. Note that $DJ_n \cdot X = (DJ)_n \cdot \mathcal{T}(\cdot, n)X$, is the instantaneous sensitivity of $J_n = J \circ \varphi_n$ to an infinitesimal perturbation to u along $X(u)$. So, the n -th summand is the ensemble sensitivity of the function J_n . In a chaotic system, the integrand exhibits exponential growth μ -a.e. The reason is that, by the definition of chaos, the norm of an infinitesimal perturbation to a μ -a.e. initial condition u , along a direction $X(u)$, asymptotically grows exponentially with time, for almost every $X(u)$. More clearly, at u μ -a.e., for almost every tangent vector

$X(u) \in T_u M$, where $\|\cdot\|$ indicates the Euclidean norm in \mathbb{R}^d ,

$$\limsup_{n \rightarrow \infty} \frac{\log \|\mathcal{T}(u, n)X(u)\|}{n} > 0.$$

For a generic J , $|D(J \circ \varphi_n) \cdot X|$ ($|\cdot|$ denotes the Euclidean norm on \mathbb{R}) shows the same asymptotic trend as $\|\mathcal{T}(u, n)X(u)\|$, that is, the integrand in Ruelle's formula, denoted as $L_n(u) := (DJ)_n \cdot \mathcal{T}(u, n) X(u)$ grows exponentially in norm, with n , at almost every u . This makes Ruelle's formula inefficient to evaluate directly, because the number of samples required to accurately obtain the n th term in the series, grows rapidly with n . It is worth noting that despite the pointwise exponential growth, since Ruelle's formula converges, i.e., the series in Eq. 5.2 converges, the ensemble averages $\langle L_n \rangle$ tend to 0 as $n \rightarrow \infty$, due to cancellations in phase space.

The pointwise exponential growth of the sensitivities is however manifest in the variance of the direct approximation of the formula, which has been shown to be computationally intractable in practical examples [69, 96]. Specifically, evaluating the series upto N , the variance of $\sum_{n=0}^N L_n$, grows exponentially with N , at almost every u . As illustrated by previous numerical results [69, 96], this leads to the least mean squared error achievable at a given computational cost, to reduce poorly with the cost, usually much worse than in a typical Monte Carlo simulation.

5.3 Main contributions

The main contribution of this chapter is a reformulation of Ruelle's formula (Eq.5.2) into a different ensemble average that provably converges like a typical Monte Carlo computation. That is, the error in the N -term ergodic average approximation of the alternative ensemble average asymptotically declines as $\mathcal{O}(1/\sqrt{N})$.

Remark 1. *Technically, the error convergence, by applying the law of the iterated logarithm [84, 174], is $\mathcal{O}(\sqrt{\log \log N}/\sqrt{N})$, but we ignore the less significant $\sqrt{\log \log N}$ throughout this chapter (see Chapters 7 and 9 for a more detailed error analysis of the S3 algorithm).*

The reformulation relies on the following result, which we also prove: the effect of an unstable perturbation on the ensemble average of an objective function is captured by its time correlation with a certain, bounded function. The formal statement is as follows.

Theorem 1. *Given an unstable covariant Lyapunov vector V^i , $1 \leq i \leq d_u$, there exist bounded scalar functions g^i such that for any $J \in C^2(M)$,*

$$\sum_{n=0}^{\infty} \langle D(J \circ \varphi_n) \cdot V^i, \mu \rangle = \sum_{n=0}^{\infty} \langle J \circ \varphi_n g^i, \mu \rangle. \quad (5.3)$$

While the proof is delayed until 5.4.4, here we mention the foremost implication of the above theorem for computation. Since the explanation in 5.2.4 applies to any unstable perturbation field, the left hand side of 5.3 does not yield a Monte Carlo computation. On the other hand, the ergodic average approximation of the right hand side does. That is, the rate of convergence would be $-1/2$, independent of the system dimension, for the right hand side. Moreover, since the right hand side enables an ergodic average approximation, a discretization of the phase space, which is impractical in high-dimensional systems, is not required.

The main result of this chapter, a computable realization of Ruelle's formula, follows as a corollary to Theorem 1, and can be stated as follows:

Corollary 1. *In uniformly hyperbolic systems, Ruelle's formula in Eq. 5.2 is equivalent to the following sum of two exponentially converging series,*

$$\left. \frac{d\langle J \rangle^s}{ds} \right|_{s=0} = \sum_{n=0}^{\infty} \langle D(J \circ \varphi_n) \cdot X^s, \mu \rangle + \sum_{n=0}^{\infty} \langle J \circ \varphi_n g, \mu \rangle, \quad (5.4)$$

where $g \in L^\infty(\mu)$, and $X = X^u + X^s$ is the decomposition of X along the unstable and stable Oseledets spaces respectively.

The regularized expression that is Eq. 5.4, is referred to as the pre-space-split sensitivity or pre-S3 formula. We briefly explain how Corollary 1 is obtained, and direct the reader to 5.4 for a complete presentation. We first split Ruelle's formula

(Eq. 5.2) into two terms by decomposing the vector field X into its unstable and stable components. The first term in Eq. 5.4 directly results from this splitting. As shown in 5.4.2, it leads to an efficient Monte Carlo computation and hence is left unchanged. The other term however, which contains the contribution from the unstable perturbation X^u , needs modification. First, it is further split into d_u terms by writing X^u in the basis of the unstable covariant Lyapunov vectors or CLVs, $V^i, 1 \leq i \leq d_u$. Assuming that each of the d_u series is well-defined, each is modified by using Theorem 1.

Before we proceed with the derivation of Eq. 5.4, we close this section by discussing the advantages offered by the new formula in Eq. 5.4. The first advantage is of course the Monte Carlo convergence of the new formula, while no such guaranteed problem-independent convergence rate can be ascribed to the direct evaluation of the original formula in Eq. 5.2. Secondly, the pre-S3 algorithm, an efficient implementation of Eq. 5.4 the details of which are deferred to a future work, only uses information obtained along trajectories, and therefore does not exhibit a direct scaling with the problem dimension. Thirdly, since Eq. 5.4 is an equivalent restatement of Ruelle's formula, the convergence of Eq. 5.4 to the true sensitivity is immediate from the convergence of Ruelle's formula [232].

5.4 Derivation of the pre-S3 formula

Our goal is to find an alternative representation of the formula 5.2 that can be easily computed. For this, we begin by splitting the parameter perturbation vector X into its stable and unstable components, along stable and unstable Oseledets spaces, denoted by E^s and E^u respectively. The motivation for splitting X becomes clear when we define the subspaces. The reader is referred to Chapter 4 of [16] for a detailed exposition on Oseledets multiplicative ergodic theorem (MET); here we use the two-sided version of the theorem for the cocycle \mathcal{T} , with the assumptions explained below. Oseledets MET gives us a direct sum decomposition (the so-called Oseledets splitting) $T_u M = E^1(u) \oplus \dots \oplus E^d(u)$, u μ -a.e., into subspaces of different asymptotic

exponential growth. The subspaces $E^i(u)$, assumed here to be one-dimensional, have the following properties:

1. covariance property: for each $i = 1, 2, \dots, d$, $\mathcal{T}(u, 1)E^i(u) = E^i(\varphi(u))$.
2. exponential growth/decay: there exist real numbers λ_i , $i = 1, \dots, d$ such that $v \in E^i(u) \neq 0 \in \mathbb{R}^d$ implies that $\lim_{n \rightarrow \pm\infty} \frac{1}{n} \log \|\mathcal{T}(u, n)v\| = \lambda_i$.

The asymptotic rates λ_i , which are called the Lyapunov exponents (LEs), are assumed to be, for simplicity, nonzero and distinct, and indexed in decreasing order, i.e., $\lambda_1 > \lambda_2 > \dots > \lambda_d$. Suppose d_u is the number of positive LEs. The unstable subspace E^u is defined as $E^u := \bigoplus_{i=1}^{d_u} E^i(u)$. As noted in 5.2.4, a chaotic system by definition has $d_u > 0$, and a nontrivial unstable subspace, consisting of nonzero vectors at μ -a.e. phase point. The stable subspace is defined as $E^s := \bigoplus_{i=d_u+1}^d E^i$.

5.4.1 Ruelle's formula split along Oseledets spaces

The unit vectors along E^i are denoted as V^i and are also called the covariant Lyapunov vectors (CLVs). Suppose X in the CLV basis (span of the V^i 's) can be expressed as $X(u) := \sum_{i=1}^d a^i(u)V^i(u)$. We write X^u and X^s to represent the decomposition of X along E^u and E^s respectively, i.e., $X^u(u) = \sum_{i=1}^{d_u} a^i(u)V^i(u)$ and $X^s(u) = \sum_{i=d_u+1}^d a^i(u)V^i(u)$. Then, we can rewrite Ruelle's formula (Eq.5.2) as,

$$\left. \frac{d\langle J \rangle^s}{ds} \right|_{s=0} = \sum_{n=0}^{\infty} \langle D(J \circ \varphi_n) \cdot X^s, \mu \rangle + \sum_{n=0}^{\infty} \langle D(J \circ \varphi_n) \cdot X^u, \mu \rangle. \quad (5.5)$$

The first term on the right hand side of the split formula (Eq.5.5) will henceforth be referred to as the stable contribution (denoted using the subscript "stable") and the second term as the unstable contribution (denoted using the subscript "unstable") to the overall sensitivity. The motivation for the split is that the stable contribution can now be computed as if the system were not chaotic, using a stable tangent equation that is developed below.

5.4.2 Derivation of the stable contribution

The stable contribution can be written as

$$\langle J, \partial_s \mu^s |_{s=0} \rangle_{\text{stable}} = \sum_{n=0}^{\infty} \langle D(J \circ \varphi_n) \cdot X^s, \mu \rangle = \sum_{n=0}^{\infty} \langle (DJ)_n \cdot \mathcal{T}(\cdot, n) X^s, \mu \rangle. \quad (5.6)$$

Now we develop a stable iterative procedure for the above expression, that satisfies our constraint of a Monte Carlo convergence, under the assumption of uniform hyperbolicity – a simplifying assumption on the dynamics that gives uniform rates of decay of perturbations along E^s and E^u forward and backward in time respectively. To wit, in a uniformly hyperbolic system, there exist constants $C, \lambda > 0$ such that $\|\mathcal{T}(u, n) X^s(u)\| \leq C e^{-\lambda n} \|X^s(u)\|$, for all $n \in \mathbb{Z}^+$ and for u μ -a.e. Such a uniform decay also applies backward in time to perturbations along E^u , i.e., $\|\mathcal{T}(u, -n) X^u(u)\| \leq C e^{-\lambda n} \|X^u(u)\|$, for all $n \in \mathbb{Z}^+$, and for u μ -a.e., with the same constants $C, \lambda > 0$. Under the assumption that φ is uniformly hyperbolic on M , and the assumption that $\|DJ\|, \|X^s\| \in L^\infty(\mu)$, at μ almost every u , $|DJ(\varphi_n(u)) \cdot \mathcal{T}(u, n) X^s(u)| \leq C e^{-\lambda n} \|DJ\|_\infty \|X^s\|_\infty$, where $\|f\|_\infty := \inf \{\alpha : \mu\{u \in M : |f(u)| > \alpha\} = 0\}$ for a scalar function $f : M \rightarrow \mathbb{R}$; when V is a vector field, $\|V\|_\infty$ is defined similarly with $|f(u)|$ replaced with $\|V(u)\|$. In other words, the L^∞ -norm of the integrand in Eq. 5.6 is exponentially decreasing with n .

5.4.3 Computation of the stable contribution

As a result of the exponentially decaying summation, truncation at a small number of terms provides a good approximation. We suggest the following method that uses a tangent equation, to compute the stable contribution in practice, since a tangent solver is usually available. We introduce the stable tangent equation, named so for using only the stable component of the perturbation but otherwise resembling a

conventional tangent equation,

$$\begin{aligned}\zeta_i^s &= D\varphi(u_{i-1})\zeta_{i-1}^s + X_i^s, \quad i = 0, 1, \dots, N-1 \\ \zeta_{-1}^s &= 0 \in \mathbb{R}^d.\end{aligned}\tag{5.7}$$

We can show that using the solutions of the above stable tangent equation, the stable contribution can be approximated as,

$$\langle J, \partial_s \mu^s |_{s=0} \rangle_{\text{stable}} \approx \frac{1}{N} \sum_{n=0}^{N-1} DJ(u_n) \cdot \zeta_n^s.\tag{5.8}$$

In 1, we show, under the assumption of uniform hyperbolicity, that the error in the above approximation decays as $\mathcal{O}(1/\sqrt{N})$.

5.4.4 The unstable contribution: an ansatz

In this section, we derive a regularized expression for the unstable contribution defined in Eq. 5.5. Denoting the components of X along the i th CLV by the scalar field a^i , we can write $X^u := \sum_{i=1}^{d_u} a^i V^i$. Thus, the unstable contribution from Eq. 5.5 can be written as,

$$\langle J, \partial_s \mu^s |_{s=0} \rangle_{\text{unstable}} = \sum_{n=0}^{\infty} \sum_{i=1}^{d_u} \langle D(J \circ \varphi_n) \cdot a^i V^i, \mu \rangle,\tag{5.9}$$

with the underlying assumption that the series in Eq. 5.9 converges for each $i \leq d_u$. We first informally motivate the ensuing derivation of a mollified expression for Eq. 5.9. The integrand can be viewed as a linear functional $l_n : E^u(u) \rightarrow \mathbb{R}$, evaluated at $V^i(u)$, and defined by $l_n(V(u)) := a^i(u)(DJ)_n \cdot \mathcal{T}(u, n)V(u)$, $V(u) \in E^u(u)$. Recall that although l_n may be bounded for a finite n , the bound is an exponentially increasing function making the evaluation of $\langle l_n(V^i) \rangle$ computationally infeasible. In particular, at μ -a.e. u , there exists an $N(u)$ such that $|l_n(V(u))| \leq \|a^i\|_{\infty} \|DJ\|_{\infty} c e^{\lambda_1 n} \|V(u)\|$, for all $n \geq N(u)$. On the other hand, due to the convergence of Ruelle's formula, the ensemble average $\langle l_n(V) \rangle$ declines asymptotically at

least when $V = V^i, i \leq d_u$. Since we ultimately want to compute $\langle l_n(V^i) \rangle$ as opposed to the pointwise values of $l_n(V^i)$, we propose the following ansatz for the unstable contribution, for some $Y^{i,n} \in E^{u*}$, the dual of E^u –

$$\langle J, \partial_s \mu^s |_{s=0} \rangle_{\text{unstable}} = \sum_{n=0}^{\infty} \sum_{i=1}^{d_u} \langle V^i \cdot Y^{i,n}, \mu \rangle. \quad (5.10)$$

In particular, we require that the vector field $Y^{i,n}$ is bounded for all n and additionally such that $\left| \sum_{i=1}^{d_u} \langle Y^{i,n} \cdot V^i \rangle \right|$ exponentially decreases with n . Then, if the central limit theorem holds for the integrand above, the heuristic expression in Eq. 5.10 leads to a desired Monte Carlo algorithm via ergodic averaging. Essentially, the ansatz chosen to satisfy our computational constraints suggests a vector field $Y^{i,n}$ that captures the overall sensitivity of J_n to perturbations along V^i . It is important that the pointwise values of $l_n(V^i)$ are not matched. The reason is, whenever for a finite n depending on u , l_n is a bounded linear functional on $E^u(u)$, the uniqueness of $Y^{i,n}(u) \in E^{u*}(u)$ from Riesz representation theorem gives $Y^{i,n}(u) = a^i(u) D J_n(u)$. As a result, $Y^{i,n}$ does not satisfy our requirements anymore and the original problem of large variances of ergodic averages has not been solved. Thus, we may require that $l_n(V^i(u))$ not equal $V^i(u) \cdot Y^{i,n}(u)$ at each u and only that $\langle l_n(V^i) \rangle = \langle V^i \cdot Y^{i,n} \rangle$, at each n .

Reformulation of the unstable contribution

In this section, beginning with the original expression in Eq. 5.5, we derive a new expression for the unstable contribution, holding Eq. 5.10 as a motivation. From Eq. 5.5, fixing an $i \leq d_u$, and isolating the n -th summand,

$$\langle a^i D(J \circ \varphi_n) \cdot V^i, \mu \rangle = \langle D(a^i J \circ \varphi_n) \cdot V^i, \mu \rangle - \langle (J \circ \varphi_n) D a^i \cdot V^i, \mu \rangle. \quad (5.11)$$

First, assuming each a^i is differentiable along E^i , Eq. 5.11 is valid. This is a rather restrictive assumption. It means that stable and unstable subspaces are differentiable functions of the phase space, a condition that is generally violated [130]. Incidentally, in the special case we consider later in this thesis, that of a one-dimensional unstable

manifold, the stable subspace is differentiable in the phase space (see e.g. Chapter 19 of [154]). In this case, since the stable component, X^s , is differentiable on the one-dimensional unstable manifold, so is X^u because X is differentiable by assumption.

The second term in Eq. 5.11 is a time correlation at time n , between the functions J and $Da^i \cdot V^i$. If both these functions are assumed to be continuous, then the correlation between them decays exponentially in time [74]. That is, the second term would approach its mean exponentially fast, for some $\gamma \in (0, 1)$:

$$|\langle (J \circ \varphi_n) Da^i \cdot V^i, \mu \rangle - \langle J, \mu \rangle \langle Da^i \cdot V^i, \mu \rangle| \sim \mathcal{O}(\gamma^n). \quad (5.12)$$

In fact, if X is assumed to be smooth, $\sum_{i=1}^{d_u} \langle Da^i \cdot V^i \rangle = 0$ (this result is proved in Theorem 3.1(b) of [232]). It then follows that $|\langle J \circ \varphi_n \sum_{i=1}^{d_u} Da^i \cdot V^i \rangle|$ exponentially decreases in n , and hence the vector field $Y_1^{i,n} := -J \circ \varphi_n Da^i$ potentially forms a part of $Y^{i,n}$. The restatement in Eq. 5.11 therefore confines the problematic derivative to the first term in Eq. 5.11. We can now focus our attention on the first term to obtain the remainder of $Y^{i,n}$. Applying measure preservation of φ on the first term, we obtain, for some $k \in \mathbb{N}$,

$$\langle D(a^i J \circ \varphi_n) \cdot V^i, \mu \rangle = \langle (D(a_k^i J_{n+k}))_k \cdot V_k^i, \mu \rangle, \quad (5.13)$$

where we have adopted the succinct notation J_k to denote the function $J \circ \varphi_k$; we neglect writing the subscript when $k = 0$. The motivation for using measure preservation forward in time becomes clear in the subsequent steps. Some intuitive reasoning can be immediately made however: $Y^{i,n}$ is a vector field that captures the ensemble average of the directional derivative of $J_n := J \circ \varphi_n$, without matching the pointwise derivatives. As a next step, we use the covariance of V^i to express the integrand in Eq. 5.13 as a linear functional on $V^i(u)$ since we want to obtain part of $Y^{i,n}(u)$.

Putting $k = 1$,

$$\langle D(a^i J_n) \cdot V^i, \mu \rangle = \langle (D(a_1^i J_{n+1}))_1 \cdot \frac{\mathcal{T}(u, 1)V^i}{z^i}, \mu \rangle = \langle D(a_1^i J_{n+1}) \cdot \frac{V^i}{z^i}, \mu \rangle, \quad (5.14)$$

where we have introduced $z^i(u) := \|\mathcal{T}(u, 1)V^i(u)\|$, and used the chain rule to go from the second expression to the third. Note that $\langle \log |z^i|, \mu \rangle = \lambda_i$, and in a uniformly hyperbolic system $\|z^i\|_\infty > e^\lambda/C$. In order to take full advantage of the downscaling offered by the z^i , we again rewrite the integrand in the following way that is valid because z^i is differentiable along E^i :

$$\langle D(a^i J_n) \cdot V^i, \mu \rangle = \langle D(a_1^i J_{n+1}/z^i) \cdot V^i, \mu \rangle - \langle J_{n+1} a_1^i D(1/z^i) \cdot V^i, \mu \rangle. \quad (5.15)$$

One advantage of rewriting is immediately clear – an infinite sum of the second term over n , to obtain the unstable contribution, is well-posed. This is because, similar to the second term in Eq. 5.11, this term is in the form of a time correlation. The other advantage in Eq. 5.15 is that it can be used iteratively to evaluate the left hand side, making the integrand asymptotically smaller in norm. We now develop such an iterative procedure by first noticing that Eq. 5.15 is valid for any bounded function J . Indeed since $\|(a_1^i J_{n+1})/(a^i z^i)\|_\infty \leq \|J\|_\infty \|1/z^i\|_\infty \leq C e^{-\lambda} \|J\|_\infty$, $(a_1^i J_{n+1})/(a^i z^i)$ is also a bounded function and thus can replace J_n on the left hand side of Eq. 5.15. Doing this replacement we obtain,

$$\langle D(a_1^i J_{n+1}/z^i) \cdot V^i, \mu \rangle = \langle D(a_2^i J_{n+2}/(z_1^i z^i)) \cdot V^i, \mu \rangle - \langle \frac{a_2^i J_{n+2}}{z_1^i} D(1/z^i) \cdot V^i, \mu \rangle. \quad (5.16)$$

Note that $(z_1^i z^i)(u) = \|\mathcal{T}(u, 2)V^i(u)\|$; for notational convenience we introduce the scalar function $y^{i,k}(u) := \|\mathcal{T}(u, k)V^i(u)\| = \prod_{j=0}^{k-1} z_j^i(u)$, $k \in \mathbb{N}$. Now Eq. 5.15 can be used as a base for recursion by substituting for the first term on its right hand side

using Eq. 5.16. Thus Eq. 5.15 becomes,

$$\langle D(a^i J_n) \cdot V^i, \mu \rangle = \langle D\left(\frac{a_2^i J_{n+2}}{y^{i,2}}\right) \cdot V^i, \mu \rangle - \sum_{k=1}^2 \langle \frac{a_k^i J_{n+k} z^i}{y^{i,k}} D(1/z^i) \cdot V^i, \mu \rangle. \quad (5.17)$$

Now the recursion can be continued by obtaining an expression for the first term on the right hand side of Eq. 5.17, by using $(a_2^i J_{n+2})/(a^i y^{i,2})$ in place of J_n and so on. We obtain the following expression in the infinite limit of applying this recursion,

$$\langle D(a^i J_n) \cdot V^i, \mu \rangle = \lim_{k \rightarrow \infty} \langle D(a_k^i J_{n+k}/y^{i,k}) \cdot V^i, \mu \rangle - \sum_{k=1}^{\infty} \langle \frac{a_k^i J_{n+k} z^i}{y^{i,k}} D(1/z^i) \cdot V^i, \mu \rangle. \quad (5.18)$$

In lemma 1, we show that the limit in the first term in Eq. 5.18 is 0. In fact, the result that is proved is that for a sequence of bounded functions f_n which goes to 0 pointwise almost everywhere, the sequence of ensemble averages of directional derivatives along the unstable directions also converges to 0. On applying measure preservation to each summand in the second term of Eq. 5.18, we obtain a series of time correlations of a function J with another bounded function, so that Eq. 5.18 becomes,

$$\langle D(a^i J_n) \cdot V^i, \mu \rangle = - \sum_{k=1}^{\infty} \langle \frac{a^i J_n z_{-k}^i}{y_{-k}^{i,k}} D(1/z_{-k}^i)_{-k} \cdot V_{-k}^i, \mu \rangle. \quad (5.19)$$

The second term in the equation above is a converging series because the L^∞ norms of the integrands are exponentially decreasing with k . More clearly, we have at μ -a.e. u that

$$\left| \frac{a^i J_n z_{-k}^i}{y_{-k}^{i,k}} (D(1/z_{-k}^i))_{-k} \cdot V_{-k}^i \right| \leq \|a^i\|_\infty \|J\|_\infty \|D(1/z^i) \cdot V^i\|_\infty \frac{z_{-k}^i}{y_{-k}^{i,k}} \leq C' e^{-\lambda(k-1)}, \quad (5.20)$$

and hence

$$\left\| \frac{a^i J_n z_{-k}^i}{y_{-k}^{i,k}} (D(1/z_{-k}^i))_{-k} \cdot V_{-k}^i \right\|_\infty \leq C' e^{-\lambda(k-1)}.$$

Thus by dominated convergence applied to the sequence $g_j, j = 1, 2, \dots$ of bounded

functions,

$$g_j := - \sum_{k=1}^j \frac{z_{-k}^i}{y_{-k}^{i,k}} (D(1/z_{-k}^i))_{-k} \cdot V_{-k}^i,$$

we obtain that $g^i := \lim_{j \rightarrow \infty} g_j^i$ is also a bounded function and that Eq. 5.18 becomes,

$$\langle D(a^i J_n) \cdot V^i, \mu \rangle = - \langle a^i J_n \left(\sum_{k=1}^{\infty} \frac{z_{-k}^i}{y_{-k}^{i,k}} (D(1/z_{-k}^i))_{-k} \cdot V_{-k}^i \right), \mu \rangle = \langle a^i J_n g^i, \mu \rangle. \quad (5.21)$$

By assumption, the summation over n of the left hand side of Eq. 5.21 converges. This implies, since the series on the right hand side must also converge, that $\lim_{n \rightarrow \infty} \langle a^i J_n g^i, \mu \rangle = 0$. On the other hand, this limit must be equal to $\langle J, \mu \rangle \langle a^i g^i, \mu \rangle$ since time correlations must decay to 0 on hyperbolic attractors. Thus, we have $\langle a^i g^i, \mu \rangle = 0$ since this is true for any bounded function J that satisfies the assumption that the series $\sum_{n=0}^{\infty} \langle D(a^i J_n) \cdot V^i, \mu \rangle$ converges. The derivation of Eq.5.21 and showing that $g^i \in L^\infty$ complete the proof of 1. Finally, note that the ansatz from section 5.4.4 is also valid. To see this, take

$$Y_2^{i,n} = - \sum_{k=1}^{\infty} \frac{a_k^i J_{n+k} z^i}{y^{i,k}} D(1/z^i),$$

and set $Y^{i,n} = Y_1^{i,n} + Y_2^{i,n}$.

5.4.5 Computation of the unstable contribution

To complete the derivation of a regularized unstable contribution, we can rewrite the first term in Eq. 5.11 by using the expression derived in Eq. 5.21. Thus, we obtain the following regularized unstable contribution,

$$\begin{aligned} \langle J, \partial_s \mu^s |_{s=0} \rangle_{\text{unstable}} &= \sum_{n=0}^{\infty} \left(\langle J_n \sum_{i=1}^{d_u} a^i g^i, \mu \rangle - \sum_{i=1}^{d_u} \langle J_n D a^i \cdot V^i, \mu \rangle \right) \\ &= \sum_{n=0}^{\infty} \langle J_n \sum_{i=1}^{d_u} (a^i g^i - D a^i \cdot V^i), \mu \rangle. \end{aligned} \quad (5.22)$$

In order to compute the unstable contribution in the form above, we resort to ergodic approximation of the ensemble average. Since we expect the time correlation

between the bounded function $g := \sum_{i=1}^{d_u} a^i g^i - Da^i \cdot V^i$ and J to decay exponentially in a uniformly hyperbolic system, the summation over n would converge (to within machine precision of the true unstable contribution) with a small number of terms, when compared to N , the trajectory length used for an ergodic average approximation of each term. Thus, the computational time for the unstable contribution is roughly equal to that for evaluating g along N points. The function g^i is naturally in the form of an iteration and thus can be obtained along a trajectory, by solving the following set of d_u scalar equations, $1 \leq i \leq d_u$, setting $\beta_{-1}^i = 0$,

$$\beta_{k+1}^i = \beta_k^i / z_k^i + D(1/z_k^i) \cdot V_k^i, \quad k = 0, 1, \dots \quad (5.23)$$

The solutions β_K^i , $K \in \mathbb{Z}^+$, approximate the scalar function g^i , asymptotically. That is, for large K , $g^i(\varphi_K(u)) \approx -\beta_K^i$. Using this approximation of g^i and a finite difference approximation of $Da^i \cdot V^i$, we can obtain the function g along a primal trajectory starting from a μ -typical phase point u . Then, the numerical approximation of the unstable contribution is the following ergodic average,

$$\frac{d\langle J \rangle}{ds} \text{ unstable} \approx \frac{1}{N} \sum_{n=0}^M \sum_{i=0}^{N-1} J(u_{n+i}) g(u_i). \quad (5.24)$$

Ignoring the numerical errors in the computation of g , 2 shows that the error in the approximation above decays as $\mathcal{O}(1/\sqrt{N})$. This completes the proof of 1.

5.5 Interpretation of the unstable contribution

In the previous section, we rewrote each term of Ruelle's formula, which represents the ensemble average of an unstable derivative, $\langle DJ_n \cdot X^u \rangle$, as a time correlation integral $\langle J_n g, \mu \rangle$ where g was a bounded distribution that we obtained through an iterative procedure. In this section, we provide physical intuition for g by relating it to the change in the SRB measure due to a perturbation along X^u .

We start with the simple case in which the SRB measure is absolutely continuous

with respect to Lebesgue measure on the whole manifold M . For the derivation of an pre-S3 formula that assumes the existence of a density on the whole manifold, see [70]. Examples of systems where this is true include expanding dynamics on compact attractors that have no stable submanifolds. In these cases, the volume element $d\mu = \rho du$, where $du = |dx_1 \cdots dx_d|$ is the standard volume element, for some smooth function $\rho : M \rightarrow \mathbb{R}^+$. Then, integration by parts of each term of the unstable contribution according to Ruelle's formula, can be performed as follows,

$$\begin{aligned} \langle D(J \circ \varphi_n) \cdot X^u, \mu \rangle &= \int_M \operatorname{div}(J \circ \varphi_n X^u) \rho \, du - \int_M J \circ \varphi_n \operatorname{div} X^u \rho \, du \\ &= \int_M \operatorname{div}(\rho J \circ \varphi_n X^u) \, du - \int_M (J \circ \varphi_n) \left(\frac{D\rho}{\rho} \cdot X^u + \operatorname{div}(X^u) \right) \rho \, du. \end{aligned} \tag{5.25}$$

By Stokes theorem, the first term in Eq. 5.25 is a boundary integral that gives the flux of the vector field X^u at the boundary of M , which is 0. Thus, in this case, the function $g \equiv -(D\rho \cdot X^u/\rho + \operatorname{div}(X^u))$. Hence, $g\rho = -\operatorname{div}(\rho X^u)$. Roughly speaking, Eq. 5.25 captures the average of the function J_n multiplied by the change in the probability distribution. Hence, this definition of ρg matches our intuition since locally, the perturbation X^u stretches the standard volume (du) by $g\rho = -\operatorname{div}(\rho X^u)$. Moreover, as derived in section 5.4, it is easy to see that $\langle g, \mu \rangle = \int g\rho \, du = 0$.

Now consider the more general case where the SRB measure is not absolutely continuous on the whole manifold. Although the derivation of Eq. 5.25 is not valid, an interpretation of the unstable contribution can be made using a similar argument. First we choose a measurable partition, say ξ , such that each partition element $\xi(u)$ that contains u lies within the local unstable manifold at u . Since conditional measures of SRB measures along unstable manifolds are absolutely continuous, (see [169, 76] for constructions of measurable partitions and disintegration of SRB measures), we can write the conditional measure of μ on $\xi(u)$ as $\rho_u(w)dw$ for some function ρ_u , where $dw = |dx_1 \cdots dx_{d_u}|$ is the standard Euclidean volume element in d_u dimensions. In coordinates, at any $w \in \xi(u)$, $X^u(w)$ can be written as $X^u(w) = \sum_{k=1}^{d_u} v^k(w)\partial_{x_k}$, for some scalar functions v^k . Using such a disintegration of the SRB measure, each

term of the unstable contribution will then have the following form, taking $a^i = 1$ for simplicity,

$$\langle D(J \circ \varphi_n) \cdot X^u, \mu \rangle = \int_{M/\xi} \int_{\xi(u)} \sum_{k=1}^{d_u} v^k(w) \frac{\partial J \circ \varphi_n}{\partial x_k} \rho_u(w) dw d\hat{\mu}, \quad (5.26)$$

where $\hat{\mu}$ is the factor measure defined as the pushforward of μ under the projection map $\pi : M \rightarrow \xi$, which maps a phase point u to $\xi(u)$, hence: $\hat{\mu} = \mu \circ \pi^{-1}$ [76]. From this point, we can treat the d_u -dimensional inner integral analogously to the previous case of the expanding map. In particular, we can apply integration by parts to the inner integral, and analogous to Eq. 5.25, we obtain,

$$\begin{aligned} \langle D(J \circ \varphi_n) \cdot X^u, \mu \rangle &= - \int_{M/\xi} \int_{\xi(u)} J \circ \varphi_n \left(\sum_{k=1}^{d_u} v^k \frac{\partial \rho_u}{\partial x_k} \right) dw d\hat{\mu} \\ &\quad - \int_{M/\xi} \int_{\xi(u)} J \circ \varphi_n \sum_{k=1}^{d_u} \frac{\partial v^k}{\partial x_k} \rho_u(w) dw d\hat{\mu}. \end{aligned} \quad (5.27)$$

Here, again we obtain an integral representing a flux term on the boundaries of $\xi(u)$, the integral over u of which is 0, due to cancellations [234]. Then, comparing with Eq. 5.22, we can see that

$$g\rho_u \equiv - \sum_{k=1}^{d_u} \left(v^k \frac{\partial \rho_u}{\partial x_k} + \frac{\partial v^k}{\partial x_k} \rho_u \right) = - \sum_{k=1}^{d_u} \frac{\partial v^k \rho_u}{\partial x_k},$$

which is again a divergence of $\rho_u X^u$ on pieces of unstable manifolds. While this provides an intuitive interpretation of g , it does not lead to a straightforward computation since the densities on the unstable manifolds, denoted ρ_u above, are unknown. This justifies resorting to an iterative procedure that we did in section 5.4, since the formula in 5.22 only makes use of known quantities computed along trajectories. The other primary motive that Eq. 5.22 fulfills is that the algorithm must not involve discretization of the phase space, but remain a Monte Carlo method of computing integrals, which have convergence rates that are independent of the dimension of the phase space.

5.6 Comments on pre-S3 computation

Revisiting the sketch of the proof in 5.3, the first term of Eq. 5.4 appears as is from the split Ruelle's formula in Eq. 5.5. Piecing together all the work carried out in 5.4.4, an exponentially converging, regularized expression for the unstable contribution, the second term of Eq. 5.5, is crystallized into Eq.5.22. Putting these two contributions together in Eq. 5.5 completes the proof of 1. Moreover, the error in the ergodic approximation of Eq. 5.22 decays as $\mathcal{O}(1/\sqrt{N})$ using an N -term ergodic average: this follows from 2. Thus, combining this result with 1, the overall pre-S3 formula has an error that decays as a typical Monte Carlo integration, as we sought.

We now briefly discuss a naïve implementation of the pre-S3 formula, postponing an efficient algorithmic implementation (see [60] for a superficial report in the case $d_u = 1$) to a future work. Using a generic initial condition u sampled according to μ on the attractor, a primal trajectory of length N , chosen large enough for convergence of ergodic averages, is obtained from the solution of Eq. 5.1. Along the primal trajectory, we use Ginelli *et al.*'s algorithm (see [114] for the algorithm and [213] for more details and a new convergence proof) to obtain $V_n^i, 0 \leq n \leq N - 1, 1 \leq i \leq d_u$. Additionally, we also apply Ginelli *et al.*'s algorithm to the adjoint cocycle (dual of \mathcal{T}) to obtain a set of adjoint CLVs, also normalized at each u , and denoted as $W_n^i, 0 \leq n \leq N - 1, 1 \leq i \leq d_u$. Note that since $E^s(u) \perp E^{u*}(u)$, $X^s \cdot W^i = (X - X^u) \cdot W^i = 0$. This fact is used in order to obtain the stable and unstable components X_n^u and X_n^s along a trajectory.

To realize the stable contribution in practice, the iterative equation referred to as the stable tangent equation (Eq. 5.7) is used, as suggested in 5.4.3. For the unstable contribution, Eq. 5.22 is computed as an ergodic average. For the computation of each g^i , Eq. 5.23 is used as suggested in 5.4.5. In the numerical examples discussed below, we use both analytical expressions and approximate finite difference calculations to obtain z^i and $D(1/z^i) \cdot V^i$, along trajectories. An algorithm for computation of derivatives of scalar functions along CLVs will be discussed in a future work, along with an adjoint (reverse-mode) algorithm for pre-S3, in the interests of serving a

high-dimensional parameter space.

Before we close this section, we comment on the uniform hyperbolicity assumption. Firstly, note that the assumption has been used to obtain the desired error convergence of both the stable and unstable contribution; the split of Ruelle's formula itself does not require uniform hyperbolicity. In particular, in the stable contribution, we used the uniform rates of contraction of stable vectors, in 1. In the unstable contribution derivation (i.e., in proving 1), and in fact in Ruelle's linear response formula itself, we use the existence of an SRB measure, which is guaranteed on a compact uniformly hyperbolic attractor. To obtain the error convergence of the unstable contribution, we used exponential decay of correlations and the CLT, which only hold on a hyperbolic attractor, for Hölder continuous functions of some positive Hölder exponent. (see [175] and section 6 of [280]). While the function J is in C^2 and hence in a Hölder class, we have only shown boundedness of g^i , but assumed exponential decay of correlations with J . However, if the two functions satisfy the finite first moment condition of Chernov (see Corollary 1.7 of [74]), the assumption of CLT and exponential decay of correlations would be valid. Moreover, besides these caveats, the assumption of uniform hyperbolicity itself could appear restrictive enough to affect the applicability of our results to high-dimensional dynamical systems encountered in practice. In this regard, it is worth mentioning that in a widely accepted hypothesis due to Gallovotti and Cohen ([108], see also [233] for more comments on this hypothesis), many fluid systems, and more generally, statistical mechanical models, behave as if they were uniformly hyperbolic. Several recent studies also provide supportive evidence, wherein numerical methods that, strictly speaking, assume some hyperbolicity for their derivation and convergence, work well in high-dimensional real-life models (see [73] for an example from climate dynamics and [208] for a turbulent fluid flow simulation).




Figure 5-1: Comparison of the sensitivities computed with pre-S3 to finite-difference for the solenoid map in Section 5.7.1. (a) J is a set of two-variable nodal basis functions along r and θ axes. (b) J_{θ} is a set of nodal basis functions along θ axis.

5.7 Numerical examples

5.7.1 Smale-Williams solenoid map

The Smale-Williams solenoid map is a classic example of low-dimensional hyperbolic dynamics. It is a three-dimensional map given by $\varphi^s(u) = [s_1 + \frac{r - s_1}{4} + \frac{\cos(\theta)}{2}, 2\theta + \frac{s_2}{4} \sin(2\pi\theta), \frac{z}{4} + \frac{\sin \theta}{2}]^T$, where $u := [r, \theta, z]^T$ in cylindrical coordinates. The attractor is a subset of the solid torus at the reference values of $s_1 = 1.4$ and $s_2 = 0$. The probability distribution on the attractor is an SRB distribution [232, 281] that has a density on the unstable manifolds. In this map, r and z directions form a basis for the stable subspace at each point (and the orthogonal θ direction forms a basis for the adjoint unstable subspace). Applying a perturbation to s_1 causes a stable perturbation, i.e., the unstable contribution is zero, since it affects only the r coordinate. On the other hand, perturbing s_2 leads to a nonzero unstable contribution. A set of nodal basis functions along r and θ is chosen to be the objective function. We use a naïve implementation of the pre-S3 algorithm presented in 5.6. In order to validate the pre-S3 computation, we compare the sensitivities $(d\langle J \rangle / ds_2)$ with finite-difference results generated using 10 billion Monte Carlo samples on the attractor. The sensitivities to




Figure 5-2: Comparison of the sensitivities of the nodal basis functions along the θ and ϕ axes to the parameter s_2 obtained for the Kuznetsov-Plykin attractor using (a) finite difference and (b) the pre-S3 algorithm.

the parameter s_2 are shown in Figure 1(a). In Figure 1(b), the objective function is a set of nodal basis functions along the θ direction. From Figures 1(a,b), we see close agreement between the sensitivities computed with (a more general version of) pre-S3 and, finite-difference results, thus validating both the stable and unstable parts of the pre-S3 algorithm.

5.7.2 Kuznetsov-Plykin map

As a second test case for pre-S3, we consider the Kuznetsov-Plykin map as defined by [160], which describes a sequence of rotations and translations on the surface of the three-dimensional unit sphere. The two parameters we choose to vary are $s_1 := \epsilon$ and $s_2 := \mu$, which are defined by [160]. The map is given by $\varphi_{n+1}^s(u) = f_{-1,-1} \circ f_{1,1}(u)$ where $u = [x_1, x_2, x_3]^T \in \mathbb{R}^3$. For the function $f_{\cdot,\cdot}$ and further details regarding the hyperbolicity of the system, the reader is referred to [160]. The probability distribution on the attractor again is again of SRB type, with the existence of a density along the unstable manifolds. We again use a naïve implementation of the pre-S3 formula to compute the sensitivities as in the case of the solenoid map in Section 5.7.1. The objective function J is a set of nodal basis functions along the θ and ϕ spherical coordinate axes. The finite-difference sensitivities were computed

with the central difference around the reference value of $s_2 = 1$ by means of 10 billion independent samples on the attractor. The results from pre-S3 agree well with finite-difference sensitivities as shown in Figure 5-2.

5.8 Conclusions

We have presented a tangent space-split sensitivity formula to compute the derivatives of statistics to system parameters in chaotic dynamical systems. The algorithm to implement the formula requires the computation of a basis for the tangent and adjoint unstable subspaces along a long trajectory. The stable contribution to the overall sensitivity can be efficiently computed by a conventional tangent/adjoint computation just as in nonchaotic systems. The unstable contribution has been rederived to be expressed as ergodic average that yields a Monte Carlo convergence. The numerical examples described in Section 5.7 satisfy the simplifying assumptions of uniform hyperbolicity and the strong differentiability of the stable and unstable subspaces, which is satisfied since both examples have one-dimensional unstable subspaces (Corollary 19.2.1 of [154]). They show close agreement with finite-difference results, serving as a proof-of-concept for the new formulation. In order to make the new formulation applicable to a high-dimensional problem, more work is needed toward an efficient implementation, particularly for the terms in Eq. 5.23. In Chapter 7, we derive a different decomposition of Ruelle's formula that does not assume the differentiability of CLVs in all directions. In fact, we do not use CLVs, which require a backward iteration to compute, but rather just an orthogonal unstable subspace, which only requires a forward pass to compute. Thus, we proceed, in Chapters 7 and 8, to a more general as well as a more computationally efficient algorithm : the S3 algorithm, although we shall restrict ourselves to uniformly hyperbolic systems with one-dimensional unstable manifolds.

5.9 Appendix

Proposition 1. *The error in an N -term ergodic approximation using the stable tangent equation 5.7, of the stable contribution, decays as $\mathcal{O}(1/\sqrt{N})$. That is,*

$$e_N := \left| \frac{d\langle J \rangle}{ds} \Big|_{\text{stable}} - \frac{1}{N} \sum_{n=0}^{N-1} DJ(u_n) \cdot \zeta_n^s \right| \leq c^s / \sqrt{N}, c^s > 0.$$

Proof. It is easy to check that $\zeta_n^s := \sum_{i=0}^n \mathcal{T}(u_i, n-i) X_i^s$ satisfies Eq. 5.7. So the approximation to the stable contribution can be written, for some $M \leq N-1$ as, where $\sum_i^j = 0$ if $i > j$,

$$\begin{aligned} \frac{1}{N} \sum_{n=0}^{N-1} \sum_{i=0}^n DJ(u_n) \cdot \mathcal{T}(u_i, n-i) X_i^s &= \frac{1}{N} \sum_{n=0}^{N-1} \sum_{i=n-M}^n DJ(u_n) \cdot \mathcal{T}(u_i, n-i) X_i^s \\ &- \frac{1}{N} \sum_{n=0}^{M-1} \sum_{i=n-M}^{-1} DJ(u_n) \cdot \mathcal{T}(u_i, n-i) X_i^s + \frac{1}{N} \sum_{n=M+1}^{N-1} \sum_{i=0}^{n-M-1} DJ(u_n) \cdot \mathcal{T}(u_i, n-i) X_i^s. \end{aligned}$$

Thus,

$$\begin{aligned} e_N &\leq \left| \frac{1}{N} \sum_{n=0}^{M-1} \sum_{i=1}^{M-n} DJ(u_n) \cdot \mathcal{T}(u_{-i}, n+i) X_{-i}^s \right| + \left| \frac{1}{N} \sum_{n=M+1}^{N-1} \sum_{i=0}^{n-M-1} DJ(u_n) \cdot \mathcal{T}(u_i, n-i) X_i^s \right| \\ &+ \left| \frac{d\langle J \rangle}{ds} \Big|_{\text{stable}} - \frac{1}{N} \sum_{n=0}^{N-1} \sum_{i=n-M}^n DJ(u_n) \cdot \mathcal{T}(u_i, n-i) X_i^s \right|. \end{aligned}$$

Under the uniform hyperbolicity assumption, we know that $\|\mathcal{T}(u, n) X^s(u)\| \leq C e^{-\lambda n}$. Moreover, we assume that $\|DJ\|, \|X^s\|$ are bounded functions. Hence, where $\gamma := \sum_{i=0}^{\infty} e^{-\lambda i}$,

$$\begin{aligned} e_N &\leq \frac{C\gamma^2 \|DJ\|_{\infty} \|X^s\|_{\infty}}{N} + \frac{C'\gamma e^{-\lambda(M+1)}(N - (M+1)) \|DJ\|_{\infty} \|X^s\|_{\infty}}{N} \\ &+ \left| \frac{d\langle J \rangle}{ds} \Big|_{\text{stable}} - \frac{1}{N} \sum_{n=0}^{N-1} \sum_{i=n-M}^n DJ(u_n) \cdot \mathcal{T}(u_i, n-i) X_i^s \right|. \end{aligned} \quad (5.28)$$

To obtain an upper bound for the third term, again we use the uniform hyperbolicity assumption. So, the integrand in the n th summand of the stable contri-

bution (Eq. 5.5) satisfies $\|(DJ)_n \cdot \mathcal{T}(\cdot, n)X^s\|_\infty \leq C \|DJ\|_\infty \|X^s\|_\infty e^{-\lambda n}$. Hence $\left\| \sum_{n=0}^M (DJ)_n \cdot \mathcal{T}(\cdot, n)X^s \right\|_\infty \leq C \|DJ\|_\infty \|X^s\|_\infty \gamma$, for any M and by dominated convergence, $\sum_{n=0}^\infty (DJ)_n \cdot \mathcal{T}(\cdot, n)X^s \in L^1(\mu)$, and the stable contribution can be written as,

$$\frac{d\langle J \rangle}{ds}_{\text{stable}} = \left\langle \sum_{i=0}^\infty (DJ)_i \cdot \mathcal{T}(\cdot, i)X^s, \mu \right\rangle.$$

Thus the third term in Eq. 5.28 has the following bound,

$$\begin{aligned} & \left| \frac{d\langle J \rangle}{ds}_{\text{stable}} - \frac{1}{N} \sum_{n=0}^{N-1} \sum_{i=n-M}^n DJ(u_n) \cdot \mathcal{T}(u_i, n-i)X_i^s \right| \\ & \leq \left| \left\langle \sum_{i=0}^M (DJ)_i \cdot \mathcal{T}(\cdot, i)X^s, \mu \right\rangle - \frac{1}{N} \sum_{n=0}^{N-1} \sum_{i=n-M}^n DJ(u_n) \cdot \mathcal{T}(u_i, n-i)X_i^s \right| \\ & + C\gamma \|DJ\|_\infty \|X^s\|_\infty e^{-\lambda M}, \end{aligned} \quad (5.29)$$

where the second term on the right hand side of Eq. 5.29 again uses uniform hyperbolicity. Applying measure preservation in each of the integrals in the first term of Eq. 5.29,

$$\begin{aligned} & \left| \frac{d\langle J \rangle}{ds}_{\text{stable}} - \frac{1}{N} \sum_{n=0}^{N-1} \sum_{i=n-M}^n DJ(u_n) \cdot \mathcal{T}(u_i, n-i)X_i^s \right| \\ & \leq \left| \left\langle \sum_{i=0}^M DJ \cdot \mathcal{T}(\varphi_{-i}(\cdot), i)X_{-i}^s, \mu \right\rangle - \frac{1}{N} \sum_{n=0}^{N-1} \sum_{i=n-M}^n DJ(u_n) \cdot \mathcal{T}(u_i, n-i)X_i^s \right| \\ & + C\gamma \|DJ\|_\infty \|X^s\|_\infty e^{-\lambda M}. \end{aligned} \quad (5.30)$$

The integrand in 5.30 is continuous on M since $\mathcal{T}(u, n) : E^s(u) \rightarrow E^s(u_n)$ is a continuous map and, $DJ(u) : T_u M \rightarrow \mathbb{R}$ and $X^s : M \rightarrow E^s$ are continuous by assumption. Then we expect that $\sum_{i=0}^N DJ \cdot \mathcal{T}(\varphi_{-i}, i)X_{-i}^s$ obeys the central limit

theorem [74]. Using this in Eq. 5.30, Eq. 5.28, gives, letting $M \rightarrow N - 1$,

$$e_N \leq \frac{C\gamma^2 \|DJ\|_\infty \|X^s\|_\infty}{N} + \frac{\text{var}[\sum_{n=0}^N DJ_n \cdot X^s]}{\sqrt{N}} + C\gamma \|DJ\|_\infty \|X^s\|_\infty e^{-\lambda(N-1)} \in \mathcal{O}(1/\sqrt{N}). \quad (5.31)$$

□

Lemma 1. *If the pointwise limit of a sequence of bounded functions $\{f_n\}_{n=0}^\infty \subset L^1(\mu)$ vanishes, i.e., $\lim_{n \rightarrow \infty} f_n(u) = 0, u \mu - \text{a.e.}$, then the sequence $\langle Df_n \cdot V, \mu \rangle$ converges to 0, when V is an unstable vector field differentiable along E^u .*

Proof. Let ξ be a measurable partition of M such that for μ -a.e. u , the element of the partition containing u , denoted $\xi(u)$, is contained in a local unstable manifold of u , i.e., $\xi(u) \subset W^u(u)$. We assume a ξ constructed according to Ledrappier-Young's Lemma 3.1.1 [169], and also use the Lyapunov-adapted coordinates introduced there. In a neighborhood of each u , let $\Phi_u : M \rightarrow [-\delta, \delta]^{d_u} \oplus [-\delta, \delta]^{d_s}$ be the adapted coordinate system such that $E^u(u), E^s(u)$ are identified with $\mathbb{R}^{d_u}, \mathbb{R}^{d_s}$ respectively. Ledrappier-Young prove the existence of a measurable function δ depending on u in order for ξ to be a measurable partition of M ; note however that in our more specific case of a uniformly hyperbolic compact attractor, we can choose a δ independently of u (see section 6.2 of [154]).

Furthermore, the image of $W^u(u)$ under Φ_u is a neighborhood of the origin in \mathbb{R}^{d_u} , i.e., the last d_s coordinates of $\Phi_u(W^u(u))$ are 0. Let the image of $\xi(u)$ under this map be $B_u \subset [-\delta, \delta]^{d_u}$. If x_1, \dots, x_{d_u} are Euclidean coordinate functions in \mathbb{R}^{d_u} , the pushforward of $V|_{\xi(u)}(w) \in E^u(w)$ through Φ_u can be expressed as $V(w) = \sum_{k=1}^{d_u} v_k(x) \partial_{x_k}, w \in \xi(u)$, and $x = \Phi_u(w)$, for differentiable functions $v_k : B_u \rightarrow \mathbb{R}$. Since $\xi(u)$ is a measurable partition, we can apply disintegration of μ on ξ , which gives for some measurable set E that $\mu(E) = \int_{M/\xi} \int_{\xi(u)} \mathbb{1}_E(w) d\mu_{\xi(u)}(w) d\hat{\mu}(\xi(u))$ [76, 281]. Here the conditional measures of μ on $\xi(u)$ are denoted as $\mu_{\xi(u)}$ and the factor measure on the quotient space M/ξ is denoted as $\hat{\mu}$. Given that μ is an SRB measure of φ , the conditional measure $\mu_{\xi(u)}$ is absolutely continuous with respect to d_u -dimensional volume measure, at μ almost every u ; let the corresponding probability

density function be denoted by $\rho_u : B_u \rightarrow \mathbb{R}^+$. Using this setup, each term of the sequence of our interest is, where $dx = |dx_1 \cdots dx_{d_u}|$ is the standard d_u -dimensional volume element, and $\tilde{f}_n := f_n \circ \Phi_u$,

$$\langle Df_n \cdot V, \mu \rangle = \int_{M/\xi} \int_{B_u} \sum_{k=1}^{d_u} v_k(x) \frac{\partial \tilde{f}_n}{\partial x_k}(x) \rho_u(x) dx d\hat{\mu}. \quad (5.32)$$

Rewriting the integrand we obtain,

$$\begin{aligned} \langle Df_n \cdot V, \mu \rangle &= \int_{M/\xi} \int_{B_u} \sum_{k=1}^{d_u} \frac{\partial}{\partial x_k} (\tilde{f}_n v_k \rho_u) dx d\hat{\mu} \\ &\quad - \int_{M/\xi} \int_{B_u} \tilde{f}_n \sum_{k=1}^{d_u} \frac{\partial (\rho_u v_k)}{\partial x_k} dx d\hat{\mu}. \end{aligned} \quad (5.33)$$

The first term goes to zero at each n due to cancellations along the boundaries of B_u at different us [234]. To see this, choose a finite cover $\cup_{l \leq r} \xi(u_l) \supset M$ and take a partition of unity supported on each $B_{u_l} = \Phi_{u_l}(\xi(u_l))$ so that the integrals on the boundaries of B_{u_l} are 0. Using dominated convergence, the second term converges to 0 as $n \rightarrow \infty$ at μ almost every u . Hence $\lim_{n \rightarrow \infty} \langle Df_n \cdot V, \mu \rangle = 0$. \square

Lemma 2. *The approximate formula Eq.5.24 for the unstable contribution has an error that decays as $\mathcal{O}(1/\sqrt{N})$: $e_{N,M} := \left| \langle J, \partial_s \mu^s \rangle_{\text{unstable}} - (1/N) \sum_{n=0}^M \sum_{i=0}^{N-1} J(u_{n+i})g(u_i) \right| \leq c^u/\sqrt{N}$, as $M \rightarrow N$, for some $c^u > 0$.*

Proof. Suppose that the central limit theorem applies to the function $\sum_{n=0}^M J \circ \varphi_n g$. Then, we can say that, for a sufficiently large N ,

$$\left| \frac{1}{N} \sum_{i=0}^{N-1} \sum_{n=0}^M J(u_{n+i})g(u_i) - \left\langle \sum_{n=0}^M J \circ \varphi_n g, \mu \right\rangle \right| \leq \frac{c_1}{\sqrt{N}}. \quad (5.34)$$

Further assuming that the decay of correlations between $J \circ \varphi_n$ and g is exponentially fast, we have, for every $n \in \mathbb{Z}^+$,

$$|\langle J \circ \varphi_n g, \mu \rangle - \langle J \rangle \langle g \rangle| \leq c_2 \gamma^n, \quad 0 \leq \gamma < 1. \quad (5.35)$$

Since we have already shown in the main text (5.4.5) that $\langle g \rangle = 0$, Eq. 5.35 gives $|\langle J \circ \varphi_n g, \mu \rangle| \leq c_2 \gamma^n$. Thus, considering the sequence of functions $h_m := \sum_{n=0}^m J \circ \varphi_n g$,

$$|\langle h_m, \mu \rangle| \leq c_2 \sum_{n=0}^m \gamma^n \leq c_2/(1 - \gamma). \quad (5.36)$$

Thus, $\{h_m\} \in L^1(\mu)$ and $\|h_m\|_1 \leq c_2/(1 - \gamma)$, and so by dominated convergence theorem, we have that, $\sum_{n=0}^{\infty} \langle J \circ \varphi_n g, \mu \rangle = \langle \sum_{n=0}^{\infty} J \circ \varphi_n g, \mu \rangle$. Then, again using Eq. 5.35, we have

$$\left| \sum_{n=0}^{\infty} \langle J \circ \varphi_n g, \mu \rangle - \left\langle \sum_{n=0}^M J \circ \varphi_n g, \mu \right\rangle \right| \leq c_2 \gamma^M / (1 - \gamma). \quad (5.37)$$

Finally,

$$e_{N,M} \leq \left| \frac{1}{N} \sum_{i=0}^{N-1} \sum_{n=0}^M J(u_{n+i}) g(u_i) - \left\langle \sum_{n=0}^M J \circ \varphi_n g, \mu \right\rangle \right| + \left| \sum_{n=0}^{\infty} \langle J \circ \varphi_n g, \mu \rangle - \left\langle \sum_{n=0}^M J \circ \varphi_n g, \mu \right\rangle \right|, \quad (5.38)$$

which using Eq. 5.34 and Eq. 5.37 gives,

$$e_{N,M} \leq c_2 \gamma^M / (1 - \gamma) + \frac{c_1}{\sqrt{N}}. \quad (5.39)$$

Taking the limit $M \rightarrow N$, this results in $e_{N,N} \in \mathcal{O}(1/\sqrt{N})$. \square

Chapter 6

Computing the curvature of one-dimensional unstable manifolds

Covariant Lyapunov vectors or CLVs span the expanding and contracting directions of perturbations along trajectories in a chaotic dynamical system. Due to efficient algorithms to compute them that only utilize trajectory information, they have been widely applied across scientific disciplines, principally for sensitivity analysis and predictions under uncertainty. In this paper, we develop a numerical method to compute the directional derivatives of the first CLV along its own direction; the norm of this derivative is also the curvature of one-dimensional unstable manifolds. Similar to the computation of CLVs, the present method for their derivatives is iterative and analogously uses the second-order derivative of the chaotic map along trajectories, in addition to the Jacobian. We validate the new method on a super-contracting Smale-Williams Solenoid attractor. We also demonstrate the algorithm on several other examples including smoothly perturbed Arnold Cat maps, and the Lorenz'63 attractor, obtaining visualizations of the curvature of each attractor. Furthermore, we reveal a fundamental connection of the derivation of the CLV self-derivative computation with an efficient computation of linear response of chaotic systems. This chapter, with minimal modification, has been published at [65], and a preprint of this chapter can be found at [63].

6.1 Introduction

Linear response refers to the linear change in the long-term or statistical behavior of a dynamical system, as a result of a small parameter perturbation. In chaotic systems, a linear response formula was developed by Ruelle [232][234], which is rigorously proved for *uniformly hyperbolic* systems, the simplest setting in which a chaotic attractor can occur. Linear response has been observed in practical chaotic systems wherein dissipative dynamics dominate [275][208][71][142]. A unique ergodic stationary physical probability distribution, known as an SRB measure [281], is achieved on uniformly hyperbolic attractors. Linear response gives us a quantitative estimate of the derivative of the SRB measure with respect to system parameters, using information only from the unperturbed system.

This statistical derivative can enable typical applications of sensitivity analysis, such as uncertainty quantification, design, optimization and control problems in chaotic systems. These applications are currently limited in chaotic systems because the computation of linear response, through Ruelle’s theoretical formula, remains a challenging problem. Some new numerical methods are being actively developed as of this writing ([62][206]; see also Chapter 7) in which Ruelle’s formula is transformed into a well-conditioned ergodic-averaging computation; other promising methods include shadowing-based methods [273][206], and approximate evaluations of Ruelle’s response using fluctuation-dissipation theorems extended to SRB-type measures [1][182][2].

In this work, we develop a numerical method for derivatives on the unstable manifold of certain quantities fundamental to linear response. These derivatives are needed for an efficient computation of a regularized version of Ruelle’s formula. Focusing on one-dimensional unstable manifolds, the proposed numerical method gives the derivative of the unstable Covariant Lyapunov Vector (CLV) [159] along its own direction. As a byproduct, we obtain the unstable derivative of the local expansion factor of the unstable CLV, and this quantity appears in the computation of linear response.

We expect the CLV self-derivatives computed in this paper, which describe the

curvature of the attractor manifold, to be applicable beyond linear response. CLVs are specific bases for tangent spaces along a trajectory, characterized by Lyapunov exponents. Ginelli *et al.*'s [114] efficient algorithm to compute CLVs has led to several applications of Lyapunov analysis in engineering, in both deterministic and stochastic chaotic systems. These applications include uncertainty quantification, data assimilation and forecasting, across a range of disciplines such as numerical weather prediction and aerospace engineering ([32], [208], [228], [142]; see [54] for a survey of applications of Lyapunov analysis).

The numerical method we develop in this work for the directional derivatives of CLVs in their respective directions is henceforth known as the *differential CLV method*. We shall refer to these derivatives as *CLV self-derivatives*. In the case of a one-dimensional unstable manifold, the CLV corresponding to the largest Lyapunov exponent is the unit tangent vector field along the unstable manifold. The norm of this CLV self-derivative is hence also the curvature of the unstable manifold.

The connection we reveal with linear response is via a byproduct of the differential CLV method: the unstable derivative of the local expansion factors of the unstable CLVs. This derivative is a key ingredient in an iterative computation of a fundamental quantity intimately connected to linear response. This quantity, which we refer to as the logarithmic *density gradient*, indicates how the SRB measure changes along unstable manifolds in the attractor. More precisely, in the case of one-dimensional unstable manifolds, which is the focus of this paper, this quantity is the unstable derivative of the logarithm of the SRB density on the unstable manifold. This connection shows one potential application of the recursive method developed in this paper: the computation of linear response in chaotic systems.

The outline of the subsequent sections is as follows. In section 6.2, we briefly summarize the theory of CLVs and establish the setting we derive our results in: uniformly hyperbolic attractors. The *differential CLV method* is derived in section 6.3; while the main steps are in section 6.3.3, notational setup and the intuition for the steps are developed in the prior subsections. We validate the method using a super-contracting Solenoid map in section 6.4.1. Further numerical experiments

demonstrating the method on the Lorenz'63 attractor, a volume-preserving perturbed Cat map, a dissipative perturbed Cat map, and the Hénon map are in sections 6.4.2, 6.4.3, 6.4.4 and 6.4.5 respectively. The implication of the method for the computation of linear response is discussed in section 6.5. We summarize our results and conclude in section 6.6.

6.2 Problem setup, definitions and review of Covariant Lyapunov Vectors

The dynamical system studied in this paper is the iterative application of a smooth (C^3) self-map $\varphi : \mathbb{M} \rightarrow \mathbb{M}$ of a domain \mathbb{M} , which is a compact subset of \mathbb{R}^m . We write φ^n to denote an n -time composition of φ ; that is, $\varphi^n = \varphi \circ \varphi^{n-1}$, $n \in \mathbb{Z}^+$, where φ^0 is the identity function on \mathbb{R}^m . The iterates under φ , or the points along orbits of the dynamical system, are represented using the following subscript notation: if $x \in \mathbb{M}$, $x_n := \varphi^n x$; x_0 is simply written as x , which we use to denote an arbitrary phase point. A similar notation is also adopted for scalar or vector-valued functions or observables. If f is an observable, $f_n := f \circ \varphi^n$. The derivative with respect to the state is denoted as d and the partial derivative operators, with respect to the Euclidean coordinate functions x_1, x_2, \dots, x_m are written as $\partial_1, \partial_2, \dots, \partial_m$, respectively. For instance, if $f : \mathbb{M} \rightarrow \mathbb{R}$ is a scalar-valued observable, the derivative df evaluated at x is given by $df(x) = [\partial_1 f(x), \dots, \partial_m f(x)]^T$. Using the notation introduced, an application of the chain rule would be as follows:

$$(df_n)^T = ((df)_n)^T d\varphi^n.$$

Finally, we assume the existence of an ergodic, physical, invariant measure for φ , known as the SRB measure and denoted μ . As a result, ergodic (Birkhoff) averages of observables in $L^1(\mu)$ converge to their expectations with respect to μ : $\lim_{N \rightarrow \infty} (1/N) \sum_{n=0}^{N-1} f_n(x) = \langle f, \mu \rangle$ for Lebesgue-a.e. $x \in \mathbb{M}$. Note that such a measure is guaranteed to exist [281] in the uniformly hyperbolic setting, which we discuss

in section 6.2.2.

6.2.1 Tangent dynamics

In order to introduce covariant Lyapunov vectors (CLVs), whose derivatives are the subject of this paper, we briefly discuss the asymptotic behavior of tangent dynamics in chaotic systems. We refer to as tangent dynamics the linear evolution of perturbations under the Jacobian matrix, $d\varphi$. The Jacobian matrix evaluated at an $x \in \mathbb{M}$ is denoted $d\varphi_x$. Denoting the tangent space at x as $T_x\mathbb{M}$, $d\varphi_x^n$ is a map from $T_x\mathbb{M}$ to $T_{x_n}\mathbb{M}$. Given a tangent vector $v_0 \in T_x\mathbb{M}$, we denote its iterate under the tangent dynamics at time n as $v_n \in T_{x_n}\mathbb{M}$. That is, $v_n = d\varphi_x^n v_0$. Intuitively, if a perturbation of norm $\mathcal{O}(\epsilon)$ is applied at x along v_0 , up to first order in ϵ , the deviation from the original orbit starting at x , after time n , is along v_n . In other words,

$$v_n = \lim_{\epsilon \rightarrow 0} \frac{\varphi^n(x + \epsilon v_0) - x_n}{\epsilon} = d\varphi_x^n v_0. \quad (6.1)$$

In practice, the above equation for the tangent dynamics is solved iteratively, along a reference orbit $\{x_n\}$, since using the chain rule, $d\varphi_x^n = d\varphi_{x_{n-1}} \cdots d\varphi_x$, and hence $v_{n+1} = d\varphi_{x_n} v_n$. A classical result in nonlinear dynamics, known as the Oseledets multiplicative ergodic theorem (OMET) [16] deals with the asymptotic behavior of v_n as $n \rightarrow \infty$, in ergodic systems. The OMET implies the following: at μ -a.e. $x \in \mathbb{M}$, the tangent space splits as a direct sum, $T_x\mathbb{M} = \bigoplus_{i=1}^p E_x^i$, $p \leq m$, where E_x^i are $d\varphi$ -invariant subspaces in the sense that $d\varphi_x E_x^i = E_{\varphi(x)}^i$. This splitting is based on the asymptotic, exponential growth/decay rates of tangent dynamics in the subspaces E_x^i . More precisely, for μ -a.e. $x \in \mathbb{M}$, if $v_0^i \in E_x^i$, its norm under the tangent dynamics grows/decays exponentially at a rate that converges to a constant. The limits

$$\lambda_{x,i} := \lim_{n \rightarrow \infty} \frac{1}{n} \log \left(\frac{\|v_n^i\|}{\|v_0^i\|} \right), \quad (6.2)$$

$1 \leq i \leq p$, are known as the Lyapunov exponents (LEs). Since φ is an ergodic map with respect to μ , the LEs are constants independent of x , for μ -a.e. x ; we denote the

LEs λ_i , in descending order as $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_p$. In our setting, φ is a chaotic map, which means that $\lambda_1 > 0$. Let d_u be the number of positive LEs, and $d_s = p - d_u$ be the number of negative LEs. Then, $E_x^u := \bigoplus_{i=1}^{d_u} E_x^i$ is called the unstable subspace of $T_x\mathbb{M}$. In other words, the unstable subspace E_x^u is the set of tangent vectors that asymptotically decay exponentially in norm under tangent dynamics backward in time; by definition, the unstable subspaces at points on a chaotic orbit are non-empty. This sensitivity to perturbations is the so-called *butterfly effect* that defines chaotic systems.

Similarly, the set of tangent vectors that asymptotically decay exponentially in norm under the tangent dynamics, make up the stable subspace, denoted $E_x^s := \bigoplus_{i=d_u+1}^p E_x^i = T_x\mathbb{M} \setminus E_x^u$. If each E^i is one-dimensional and $p = m$, the covariant Lyapunov vectors or CLVs, denoted as V^i in this paper, are unit vector fields along E^i . That is, CLVs satisfy the following properties at μ -a.e. x :

- The covariance property:

$$d\varphi_x V_x^i \in E_{x_1}^i. \quad (6.3)$$

Since by definition V_x^i is a unit vector, we introduce a scalar function $z_{\cdot,i} : \mathbb{M} \rightarrow \mathbb{R}^+$ defined as $z_{x,i} := \|d\varphi_x V_x^i\|$, to indicate the local stretching or contraction factor of the i th CLV. Hence, the covariance property of the i th CLV can be expressed as

$$d\varphi_x V_{x,i} = z_{x,i} V_{x_1}^i. \quad (6.4)$$

- The i th CLV grows/decays asymptotically on an exponential scale, at the rate λ_i , and, in addition, is invariant under time-reversal:

$$\lambda_i := \lim_{n \rightarrow \pm\infty} \frac{1}{n} \log \|d\varphi_x^n V_x^i\|. \quad (6.5)$$

6.2.2 Uniform hyperbolicity

We consider an idealized class of chaotic systems known as uniformly hyperbolic systems, which are characterized by uniform expansions and contractions of tangent vectors. In uniformly hyperbolic systems, there exist constants $c > 0$ and $\lambda \in (0, 1)$ such that, at every point $x \in \mathbb{M}$, i) every stable tangent vector $v \in E_x^s$ satisfies: $\|d\varphi_x^n v\| \leq c \lambda^n \|v\|$, and ii) every unstable tangent vector $v \in E_x^u$ satisfies: $\|d\varphi_x^{-n} v\| \leq c \lambda^n \|v\|$, for all $n \in \mathbb{N}$. As a result, in these systems, there exists an upper (lower) bound that is independent of the base point x , on the slowest contracting (stretching) factors among $z_{x,i}$. In particular, defining $C := c\lambda$, we have $z_{x,i} \geq (1/C)$, $1 \leq i \leq d_u$ and $z_{x,i} \leq C$, $d_u + 1 \leq i \leq d$. From the definition of the LEs (Eq. 6.5), it is also clear that they are the ergodic (Birkhoff) averages of the stretching/contraction factors:

$$\langle z_{\cdot,i}, \mu \rangle := \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{n=0}^{N-1} \log z_{x_n,i} = \lambda_i, \quad x \in \mathbb{M} \quad \mu - \text{a.e.} \quad (6.6)$$

6.2.3 Examples

A simple example of a uniformly hyperbolic system is Arnold's Cat map, a smooth self-map of the surface of the torus ($\mathbb{T}^2 \equiv \mathbb{R}^2/\mathbb{Z}^2$):

$$\varphi([x_1, x_2]^T) = \begin{bmatrix} 2 & 1 \\ 1 & 1 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} \pmod{1}. \quad (6.7)$$

This is a linear hyperbolic system, i.e., the Jacobian matrix of the map is a constant in phase space and has eigenvalues other than 1. In this simple example, the CLVs and the stretching/contracting factors, are also independent of the phase point. The logarithm of the eigenvalues of the constant Jacobian matrix, are the LEs of this map: $\lambda_1 = \log |(3 + \sqrt{5})/2|$ and $\lambda_2 = \log |(3 - \sqrt{5})/2|$. It is also clear that $E^1 = E^u$ and $E^2 = E^s$ are one-dimensional subspaces spanned by V^1 and V^2 , the eigenvectors of the Jacobian matrix at eigenvalues of e^{λ_1} and e^{λ_2} respectively. Moreover, z_1 and z_2 are also constant on $\mathbb{R}^2/\mathbb{Z}^2$: $z_1 = e^{\lambda_1}$, and $z_2 = e^{\lambda_2}$. Further, the SRB measure for this map is the Lebesgue measure on $\mathbb{R}^2/\mathbb{Z}^2$.

Since the Jacobian matrix is symmetric, the CLVs V^1 and V^2 are everywhere orthogonal to each other, but it is worth noting that this is a special case. In a generic uniformly hyperbolic system, it is only true that the angle between the CLVs is uniformly bounded away from zero. The perturbed Cat maps treated later have additive perturbations to the Cat map above that are smooth functions on the torus. Two types of smooth perturbations are considered later, both designed to produce non-uniform behavior of the CLVs. Both perturbed Cat maps are still uniformly hyperbolic, and differ in whether or not the resulting maps are area-preserving, in order to represent the two distinct cases of conservative (symplectic) and dissipative chaos.

6.2.4 Lack of differentiability of E^u and E^s

On hyperbolic sets, it is known that E^u and E^s are Hölder-continuous functions of phase space, in a sense clarified in Appendix section 6.7.1. When the Hölder exponent β , from Appendix section 6.7.1, equals 1, we have Lipschitz continuity, but this is indeed rare. Several examples (see [130] and references therein) have been constructed in which β is made to be arbitrarily small at almost all phase points, even in C^∞ maps. In rare cases, E^u and E^s are continuously differentiable when a certain bunching condition ([130], or section 19.1 of [154]) is satisfied by the LEs.

Revisiting the examples, the perturbed Cat maps discussed above belong to the rare category of maps with continuously differentiable stable/unstable subspaces. In fact, it can be shown that all uniformly hyperbolic maps on compact sets of dimension 2 belong to this category (see Corollary 19.1.11 of [154]). While it would be typical of a higher-dimensional map, even when uniformly hyperbolic, to show non-smoothness of the stable and unstable subspaces, we have chosen to work with two-dimensional examples in this paper for easy visualization of the subspaces, which are lines in these maps.

6.2.5 Derivatives of CLVs in their own directions

While the CLVs may lack differentiability on \mathbb{M} , they have directional derivatives in their own directions. In fact, it can be shown that these directional derivatives, which we refer to here as CLV self-derivatives, are themselves Hölder continuous with the same exponent β (see Remark in the proof of Theorem 19.1.6 of [154]). To wit, in two-dimensional uniformly hyperbolic systems, examples of which are considered in this paper, both partial derivatives (along coordinate directions) of the CLVs exist, and hence the CLVs have directional derivatives in all directions. The purpose of this paper, however, is to numerically compute directional derivatives of CLVs along their respective directions in a general uniformly hyperbolic system, regardless of their differentiability in phase space. Thus, we compute the CLV self-derivatives, without using the partial derivatives along coordinate directions, which may not exist. The CLV self-derivatives are denoted by $W_x^i \in T_x T_x \mathbb{M} \equiv \mathbb{R}^m$. They are defined using curves $\mathcal{C}_{x,i} : [-\epsilon_x, \epsilon_x] \rightarrow \mathbb{M}$ with the properties: i) $\mathcal{C}_{x,i}(0) = x$, ii) $\mathcal{C}'_{x,i}(t) = V^i(\mathcal{C}_{x,i}(t))$, $\forall t \in [-\epsilon_x, \epsilon_x]$, as

$$W_x^i := \lim_{t \rightarrow 0} \frac{V^i_{\mathcal{C}_{x,i}(t)} - V_x^i}{t}. \quad (6.8)$$

For example, in the case of a 1-dimensional unstable manifold, the curve $\mathcal{C}_{x,1}$, coincides with a local unstable manifold at x . Further discussion on the definition of W^i based on these curves, is postponed until section 6.3.1. Here we explain the existence of these curves. The vector fields V^i , $1 \leq i \leq d_u$ are infinitely smooth on an open set in a local unstable manifold, and likewise, V^i , for $d_u + 1 \leq i \leq d$ are infinitely smooth on an open set in a local stable manifold. As a result, due to the existence and uniqueness theorem, the flow of vector field V^i , denoted by the curve $\mathcal{C}_{\cdot,i}$ exists and is uniquely defined, for some $\epsilon. > 0$, justifying the definition in Eq. 6.8.

Given $T_x \mathbb{M}, T_x T_x \mathbb{M} \equiv \mathbb{R}^m$, we write all vectors in these spaces in Euclidean coordinates. The output of the numerical method to be developed, W^i , are d -dimensional vector fields consisting of component-wise directional derivatives of V^i .

6.2.6 Computations along trajectories

Before we delve into the differential CLV method, we note that W^i , being self-derivatives of CLVs, are naturally defined along trajectories, just like the CLVs. Thus, we seek a trajectory-based iterative procedure to compute them. We assume as input to the method the map, its Jacobian and second-order derivative, all computed along a long, μ -typical trajectory. The CLVs that need to be differentiated are also assumed as input, along the trajectory. To compute the CLVs, a standard algorithm such as Ginelli *et al.*'s algorithm [114] can be used. This is an iterative procedure involving repeated QR factorizations of nearby subspaces to the one that is spanned by the required CLVs. For Ginelli *et al.*'s algorithm, the reader is referred to [114] and [213] for its convergence with respect to trajectory length; for other algorithms that involve LU factorizations instead of QR, we refer to [159].

Besides using the computed CLVs as input, the differential CLV method we develop here for W^i does not follow Ginelli *et al.*'s or other algorithms for the computation of CLVs, primarily because the vector fields W^i do not satisfy the covariance property. But the method resembles the latter algorithms in being iterative and trajectory-based. One advantage of trajectory-based computation is that we exploit for fast convergence (this aspect again being similar to the CLV computation algorithms) the hyperbolic splitting of the tangent space. This will be clear at the end of the next section in which we give a step-by-step derivation.

6.3 An algorithm to compute the directional derivatives of CLVs in their own directions

In this section, we derive a numerical method to determine the quantity of interest, W^i , which is defined in Eq. 6.8. In particular, fixing a reference trajectory x, x_1, \dots , we develop an iterative scheme that converges asymptotically to vectors $W_n^i := W_{x_n}^i$, under certain conditions (Appendix 6.8), starting from an arbitrary guess for $W_0^i := W_x^i \in \mathbb{R}^m$. The derivation results in the following iteration, valid for $1 \leq i \leq d_u, n \in$

\mathbb{Z}^+ , and guaranteed to converge when $i = 1$:

$$W_{n+1}^i = \left(I - V_{n+1}^i (V_{n+1}^i)^T \right) \frac{d^2 \varphi(x_n) : V_n^i V_n^i + d\varphi(x_n) W_n^i}{z_{n,i}^2}. \quad (6.9)$$

The iteration mainly uses the chain rule and the covariance property of V^i , in a convenient set of coordinate systems centered along each μ -typical trajectory. These trajectory-based coordinates help us uncover each term on the right hand side of Eq. 6.9.

6.3.1 Change of coordinates and associated notation

Fix a μ -typical point $x \in \mathbb{M}$, and consider again the curves $\mathcal{C}_{x,i}$, $1 \leq i \leq d$, which were introduced to define W^i in Eq. 6.8. To reiterate, the curves $\mathcal{C}_{x,i} : [-\epsilon_x, \epsilon_x] \rightarrow \mathbb{M}$ are such that i) $\mathcal{C}_{x,i}(0) = x$ and ii) $\mathcal{C}'_{x,i}(t) = V^i(\mathcal{C}_{x,i}(t))$, for all $t \in [-\epsilon_x, \epsilon_x]$. There exists a measurable function $x \rightarrow \epsilon_x$ that defines the extent of the curves so that such a coordinate change, from $[-\epsilon_x, \epsilon_x]^m$ to a neighborhood of x , exists and is additionally differentiable. This follows from an assertion proved in standard stable-unstable manifold theory: a closed ϵ_x Euclidean ball around the origin in \mathbb{R}^{d_u} (\mathbb{R}^{d_s}) has an embedding into a local unstable (stable) manifold at x . These pointwise coordinate systems are referred to as Lyapunov charts or adapted coordinates in the theoretical literature ([154] Ch. 6, [169]).

In writing Eq. 6.8, we made a particular choice of adapted coordinates. We chose coordinate functions that are adapted specifically to the CLVs, as opposed to any other basis of $T_x \mathbb{M}$, in the following sense. At each x , the image of the i th Euclidean basis vector e_i , under the differential of the coordinate change, is V^i . More intuitively, we have chosen adapted coordinates such that the i th Euclidean coordinate axis corresponds, under these coordinate changes, to points that are perturbations along V^i . Thus, our quantity of interest, can be written, by definition of CLV-adapted coordinates, as

$$W_x^i = \frac{d}{dt} (V^i \circ \mathcal{C}_{x,i})(0). \quad (6.10)$$

6.3.2 The map in adapted coordinates

Now we introduce the transformation induced by $\varphi : \mathbb{M} \rightarrow \mathbb{M}$ on the CLV-adapted coordinates on \mathbb{R}^m . To do that, we fix an $i \leq d_u$ and focus on the relationship between the curves $\mathcal{C}_{x_1,i} : [-\epsilon_{x_1}, \epsilon_{x_1}] \rightarrow \mathbb{M}$ and $\varphi \circ \mathcal{C}_{x,i} : [-\epsilon_x, \epsilon_x] \rightarrow \mathbb{M}$. Define $f_{x,i} := (\mathcal{C}_{x_1,i})^{-1} \circ \varphi \circ \mathcal{C}_{x,i}$, noting that this definition makes sense at a point $t \in [-\epsilon_x, \epsilon_x]$ whenever $\varphi(\mathcal{C}_{x,i}(t))$ lies in the image of $\mathcal{C}_{x_1,i}$. The function, $x \rightarrow \epsilon_x$, which determines the size of the local unstable manifold at each x , can be chosen such that orbits of

$$f_{x,i}^{-n} := f_{x_n,i}^{-1} \circ \cdots \circ f_{x_1,i}^{-1} \circ f_{x,i}^{-1}, n \in \mathbb{Z}^+ \quad (6.11)$$

are well-defined at almost every x , for $1 \leq i \leq d_u$, within local unstable manifolds centered along the backward φ -orbit. Clearly, 0 is a fixed point of $f_{x,i}^n$ for all $n \in \mathbb{Z}$, and corresponds to the φ -orbit $\cdots, x_{-1}, x, x_1, x_2, \cdots, .$. Intuitively, if an orbit of $f_{x,i}$ excluding the fixed point, say $\{t_n := f_{x,i}^n(t)\}$, exists, it means that $\mathcal{C}_{x_n,i}(t_n)$ lies in sufficiently small local unstable manifolds of x_n , at each n . The sizes of the local unstable manifolds can be controlled in order for such orbits to be well-defined (in particular, see Lemma 2.2.2 of [169]).

To summarize, we make a specific choice of $x \rightarrow \epsilon_x$ such that the curves $\mathcal{C}_{x_n,i}$ at each n lie inside a local unstable manifold at x_n , and are tangent to $V_n^i := V_{x_n}^i$. This allows us to obtain expressions for the derivative of a CLV V_{n+1}^i with respect to V_n^i , which will in turn enter into the computation of W^i . In particular, using CLV-adapted coordinates, a suitable map $x \rightarrow \epsilon_x$, as described above, and the definition of $f_{x,i}$,

$$(df_{x,i}/dt)(0) = z_{x,i}. \quad (6.12)$$

Now we usefully relate the iterates through φ of the differential operator on \mathbb{M} corresponding to the vector field V^i , and its analog on \mathbb{R} : d/dt , along the trajectory lying in the unstable manifold of x_n . In particular, for the function V^i , when combined

with Eq. 6.10, and Eq. 6.12,

$$\begin{aligned} \frac{d(V^i \circ \varphi \circ \mathcal{C}_{x,i})}{dt}(0) &= \frac{d(V^i \circ \mathcal{C}_{x_1,i} \circ f_{x,i})}{dt}(0) \\ &= z_{x,i} W_{x_1}^i. \end{aligned} \quad (6.13)$$

6.3.3 Computation of unstable CLV self-derivatives

Starting from Eq. 6.13, and by definition of CLVs (Eq. 6.4)

$$W_{x_1}^i = \frac{1}{z_{x,i}} \frac{d}{dt} \left(\frac{d\varphi_{\mathcal{C}_{x,i}} V_{\mathcal{C}_{x,i}}^i}{z_{\mathcal{C}_{x,i},i}} \right) (0) \quad (6.14)$$

$$\begin{aligned} &= \frac{1}{z_{x,i}^2} \frac{d}{dt} (d\varphi_{\mathcal{C}_{x,i}}) (0) V_x^i + \frac{1}{z_{x,i}^2} d\varphi_x \frac{d}{dt} (V_{\mathcal{C}_{x,i}}^i) (0) \\ &+ V_{x_1}^i \frac{d}{dt} \left(\frac{1}{z_{\mathcal{C}_{x,i},i}} \right) (0) \end{aligned} \quad (6.15)$$

By Eq. 6.10, we can write the second term above as $(1/z_{x,i})^2 d\varphi_x W_x^i$. The first term can be written using the chain rule in terms of the $m \times m \times m$ second-order derivative of φ , which is a bilinear form denoted as $d^2\varphi$. Let the elements of the second-order derivative of the map be indexed such that $d^2\varphi[i, j, k] = \partial_k \partial_j \varphi_i$, and let $d^2\varphi : b$ indicate the $m \times m$ matrix resulting from taking the dot product of the last axis of $d^2\varphi$ and the vector b . Then, Eq. 6.15 becomes

$$\begin{aligned} W_{x_1}^i &= \frac{1}{z_{x,i}^2} d^2\varphi_x : V_x^i V_x^i + \frac{1}{z_{x,i}^2} d\varphi_x W_x^i \\ &+ V_{x_1}^i \frac{d}{dt} \left(\frac{1}{z_{\mathcal{C}_{x,i},i}} \right) (0) \end{aligned} \quad (6.16)$$

6.3.4 The differential CLV method: iterative orthogonal projections

The differentiation in the third term in Eq. 6.16, carried out explicitly gives,

$$\begin{aligned}
\frac{d}{dt} \left(\frac{1}{z_{\mathcal{C}_{x,i,i}}} \right) (0) &= -\frac{1}{2z_{\mathcal{C}_{x,i,i}}^3} \frac{d}{dt} \left((d\varphi_{\mathcal{C}_{x,i}} V_{\mathcal{C}_{x,i}}^i)^T d\varphi_{\mathcal{C}_{x,i}} V_{\mathcal{C}_{x,i}}^i \right) (0) \\
&= -\frac{(d\varphi_x V_x^i)^T}{z_{x,i}^3} \left(d^2\varphi_x : V_x^i V_x^i + d\varphi_x W_x^i \right) \\
&= -\frac{(V_{x_1}^i)^T}{z_{x,i}^2} \left(d^2\varphi : V_x^i V_x^i + d\varphi_x W_x^i \right). \tag{6.17}
\end{aligned}$$

Substituting Eq. 6.17 into Eq. 6.16, we see that Eq. 6.16 simply projects out the component along the $V_{x_1}^i$ direction. That is,

$$W_{x_1}^i = \left(I - V_{x_1}^i (V_{x_1}^i)^T \right) \left(\frac{d^2\varphi_x : V_x^i V_x^i + d\varphi_x W_x^i}{z_{x,i}^2} \right), \tag{6.18}$$

where I is the $m \times m$ Identity matrix. That is, the CLV self-derivatives are orthogonal to the corresponding CLVs. Before the orthogonal projection, the component along V^i is given by Eq. 6.17, which indicates the change of (the reciprocal of) the expansion factor $z_{\cdot,i}$ along V^i . This is a fundamental quantity that influences the unstable derivative of the conditional density of the SRB measure on the unstable manifold, and will be denoted

$$\alpha_{x,i} := \frac{d}{dt} \left(\frac{1}{z_{\mathcal{C}_{x,i,i}}} \right) (0).$$

We will henceforth refer to Eq. 6.17 as the *differential expansion* equation, and see its connection to linear response in section 6.5.

Now, Eq. 6.18 can be marched forward in time recursively by replacing $W_{x_1}^i$ with $W_{x_2}^i$, and W_x^i with $W_{x_1}^i$. Fixing an x , we use the subscript notation, e.g. $W_n^i := W_{x_n}^i$, and start from a random initial vector $\in \mathbb{R}^m$ as a guess for $W_0^i := W_x^i$. The following

iteration is proposed as the differential CLV method to obtain W_n^i , $n \in \mathbb{Z}^+$, $1 \leq i \leq d_u$

$$W_{n+1}^i = \left(I - V_{n+1}^i (V_{n+1}^i)^T \right) \left(\frac{(d^2\varphi)_n : V_n^i V_n^i + (d\varphi)_n W_n^i}{z_{n,i}^2} \right). \quad (6.19)$$

In Appendix section 6.8, we show that the above equation always converges asymptotically at an exponential rate when $i = 1$. For other indices $1 < i \leq d_u$, the convergence is under certain conditions on the LEs. Thus, from here on, we restrict ourselves to chaotic attractors with one-dimensional unstable manifolds, where we know the differential CLV method converges asymptotically. It is worth noting, in the context of linear response computation (section 6.5), that the differential expansion/contraction equation (Eq. 6.17) is also effectively time-evolved in order to compute the projection term in the differential CLV method (Eq. 6.19). Thus, we obtain the scalars $\alpha_{n,i}$ along a trajectory as a byproduct.

Remark 2. *Note that the entire procedure above was derived for the unstable CLV self-derivatives. For the stable ones, we must apply the same procedure with time reversal since the stable and unstable CLVs are the same, except their roles are exchanged upon time reversal. That is, when $d_u + 1 \leq i \leq m$, we must apply the above iterative procedure (Eq. 6.19) by replacing φ with the inverse map, φ^{-1} . Analogously, our numerical procedure converges when using φ^{-1} , as shown in Appendix section 6.8, for $i = m$ – for the self-derivative of the most stable CLV.*

Remark 3. *The differential CLV method may be viewed as a fast computation, exploiting dynamical information, of the extrinsic curvature of lower-dimensional smooth manifolds (in this chapter, we focus on local one-dimensional unstable manifolds) embedded in a Euclidean space. Related are computations of derivatives on Grassmanian manifolds consisting of lower-dimensional linear subspaces of Euclidean spaces that appear to be useful for geometric insight into dynamical systems problems arising in different contexts, such as in deep neural networks and reduced-order stochastic models [220, 97].*

6.4 Numerical results implementing the differential CLV method

In this section, we implement the differential CLV algorithm discussed in the previous section to several examples of low-dimensional chaotic attractors, some of which were introduced in section 6.2. In every example, the unstable subspace is one-dimensional (a line) and numerical estimates of W^1 are shown. The Python code for the implementation, along with the files needed to generate the plots in this section, can be found at [57].

6.4.1 Validation against analytical curvature of the Solenoid map

The Smale-Williams Solenoid map produces a well-known example of a uniformly hyperbolic attractor that is contained in a solid torus. We consider a two-parameter Solenoid map, which in cylindrical coordinates, is written as follows:

$$\varphi([r, t, z]^T) = \begin{bmatrix} s_0 + (r - s_0)/s_1 + (\cos t)/2 \\ 2t \\ z/s_1 + (\sin t)/2 \end{bmatrix}. \quad (6.20)$$

Clearly, the parameter s_1 is a contraction factor along the \hat{r} and \hat{z} directions. In the limit $s_1 \rightarrow \infty$, the attractor of the map, henceforth referred to as the super-contracting Solenoid attractor, becomes a space curve. It is described by the following curve parameterized by the coordinate t , expressed in Cartesian coordinates:

$$\gamma(t) := \begin{bmatrix} x_{1,n+1} \\ x_{2,n+1} \\ x_{3,n+1} \end{bmatrix} = \begin{bmatrix} \left(s_0 + \frac{\cos t}{2} \right) \cos 2t \\ \left(s_0 + \frac{\cos t}{2} \right) \sin 2t \\ \frac{\sin t}{2} \end{bmatrix}, \quad (6.21)$$




Figure 6-1: Comparison of the x_1, x_2, x_3 components of V^1 computed analytically (orange circles) and numerically (blue crosses), for the super-contracting Solenoid map




Figure 6-2: Comparison of the x_1, x_2, x_3 components of W^1 computed analytically (orange circles) and numerically (blue crosses), for the super-contracting Solenoid map.

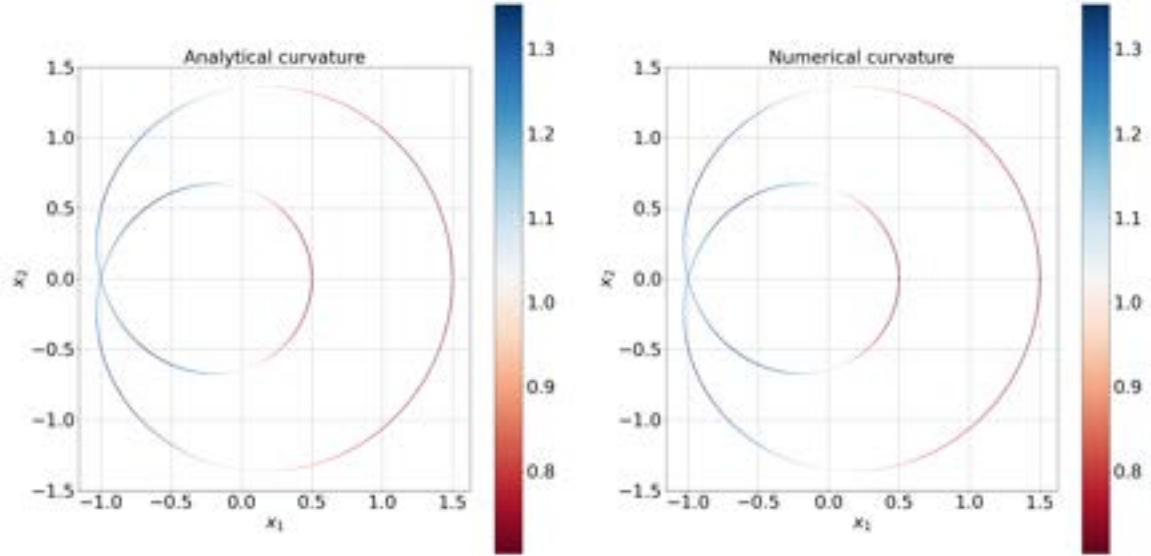


Figure 6-3: The vector field V^1 is shown for the Solenoid map. The color represents $\|W^1\|$, which is the curvature of the attractor.

where $t = \arctan(x_{2,n}/x_{1,n})$. As an aside, note that in the \hat{t} direction, the map is simply a linear expanding map, and hence the \hat{t} component of the state vector has a uniform probability distribution in $[0, 2\pi)$. We fix s_0 at 1 throughout. The one-dimensional unstable manifold is given by the curve $\gamma(t)$ defined in Eq. 6.21. Then, the tangent vector field to the curve, $\gamma'(t)$, must be along $V^1(\gamma(t))$. This is verified numerically in Figure 6-1, where the numerically computed vector field V^1 agrees closely with the unit tangent vector field $\gamma'(t)/\|\gamma'(t)\|$: in each of the subfigures, the components of the two vector fields lie superimposed on each other. Consequently, the acceleration along the curve $\gamma(t)$, $\partial_{\gamma'(t)}\gamma'(t)$, must be in the direction of $W^1(\gamma(t))$. In particular, the acceleration in the direction of the unit tangent vector, $\partial_{\gamma'(t)/\|\gamma'(t)\|}(\gamma'(t)/\|\gamma'(t)\|)$, must match $W^1(\gamma(t))$. This is also clearly seen numerically. In Figure 6-2, each component of the two vector fields $\partial_{\gamma'/\|\gamma'\|}(\gamma'/\|\gamma'\|)$, computed analytically, and W^1 , computed numerically using Eq. 6.19, are seen to coincide. Thus, the norms of the two vector fields are of course in close agreement as well, as can be seen in Figure 6-3. Both the analytically computed norm $\|\partial_{\gamma'/\|\gamma'\|}(\gamma'/\|\gamma'\|)\|$, and the numerically computed $\|W^1\|$ are shown as a colormap on the vector field $V^1 = \gamma'(t)/\|\gamma'(t)\|$. The plots in Figure 6-3 are, a fortiori, a visualization of the




Figure 6-4: Orbit points of the Lorenz system shown on the x_1 - x_3 plane, at $T_1 = 18$ (left) and at $T_2 = 20$ (right), colored according to the distance from their centroid normalized by the centroid z -coordinate. The initial conditions were 10001 equispaced points on the short line segment joining $(-0.01, 0, 1)$ and $(0.01, 0, 1)$.

curvature of the one-dimensional unstable manifold $\gamma(t)$. The final results of the analytical curvature calculations are provided in Appendix section 6.7.2.

6.4.2 Numerical verification of the curvature of the Lorenz attractor

Next we consider the well-known Lorenz'63 system [179], given by the following system of ODEs:

$$\frac{d}{dt} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = F([x_1, x_2, x_3]^T) := \begin{bmatrix} 10(x_2 - x_1) \\ x_1(28 - x_3) - x_2 \\ x_1x_2 - \frac{8x_3}{3} \end{bmatrix}. \quad (6.22)$$

The map φ is defined here to be a time-discretized form of the above system of ODEs. In particular, we use a second-order Runge-Kutta scheme with a time step of $\delta t = 0.01$. The map $\varphi(x) = x_1$, is the time-integrated solution after time δt , starting from $x := [x_1, x_2, x_3]^T \in \mathbb{R}^3$. The Lorenz'63 map defined this way has the following Lyapunov exponents: $\lambda_1 \approx 0.9$, $\lambda_2 \approx 0$ and $\lambda_3 \approx -14.6$. The unstable manifold,

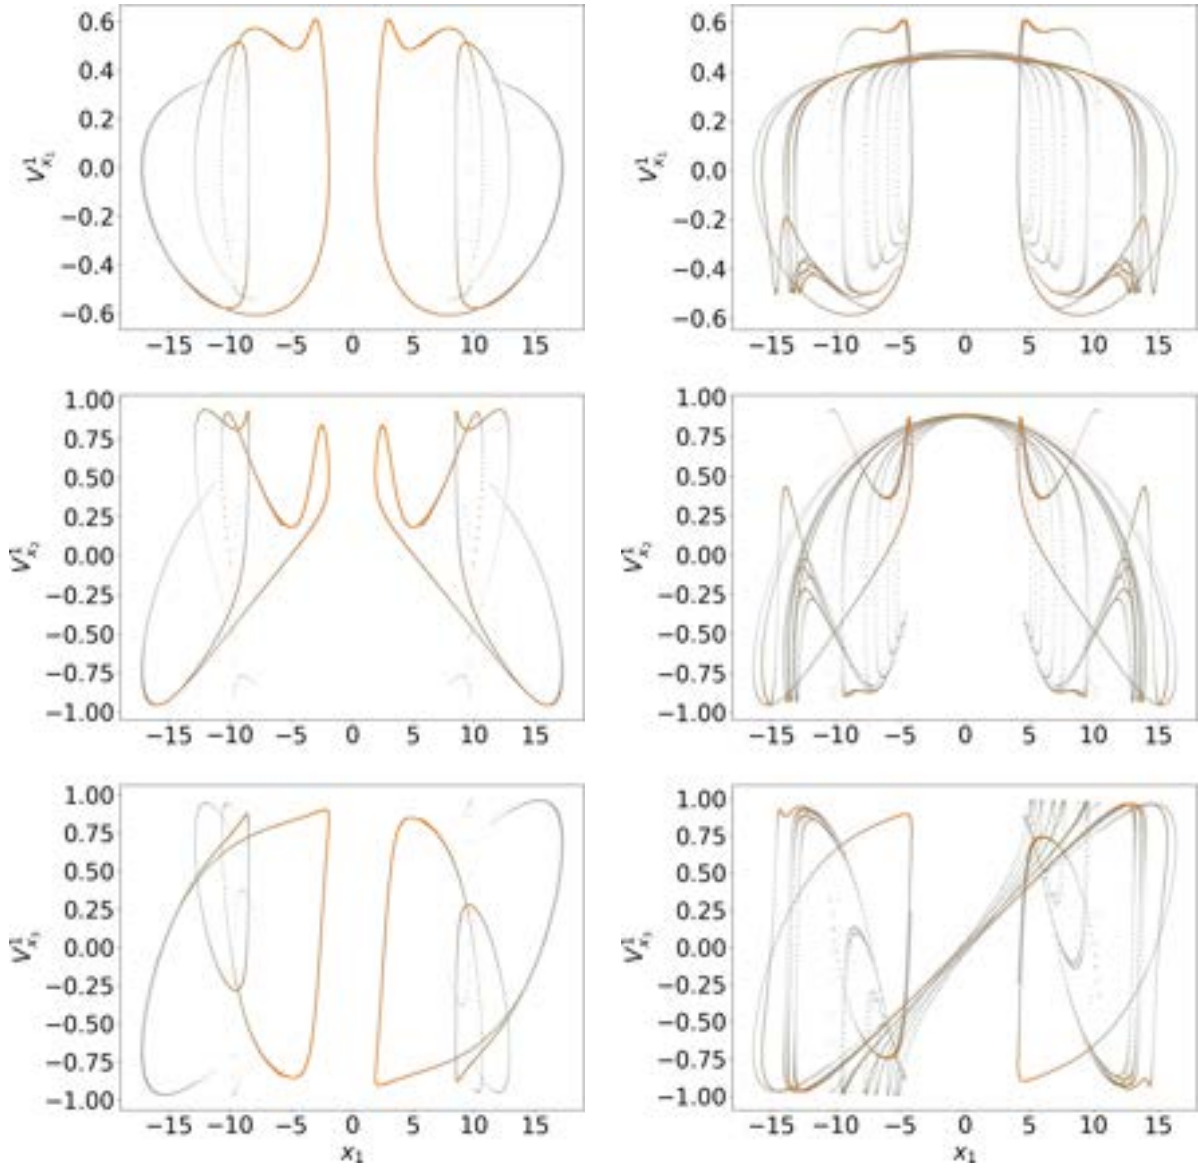


Figure 6-5: Comparison between V^1 from an iteration of the tangent dynamics (shown in orange) and V^1 from finite difference of the primal trajectories (in blue). The first column shows the components of V^1 at time $T_1 = 18$ and the second column at $T_2 = 20$. The first, second and third rows show the x_1 , x_2 , x_3 components of V^1 respectively.

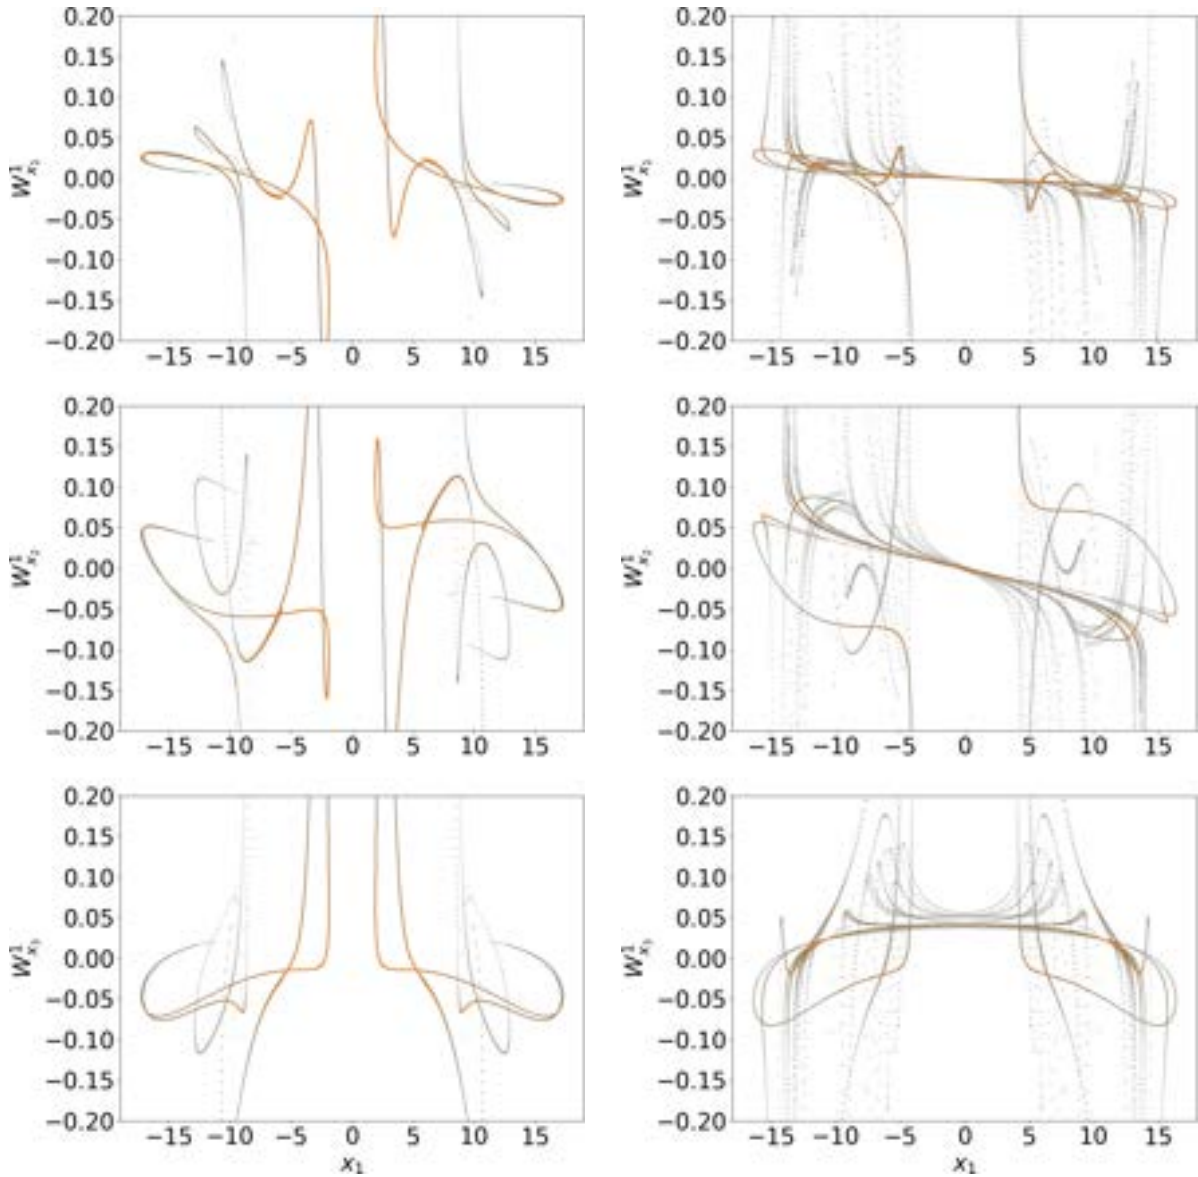


Figure 6-6: Comparison between W^1 from the differential CLV method (shown in orange) and W^1 from finite difference (in gray). The first column shows the components of W^1 at time $T_1 = 18$ and the second column at $T_2 = 20$. The first, second and third rows show the x_1 , x_2 , x_3 components of W^1 respectively.

which is tangent to the CLV corresponding to λ_1 , is one-dimensional. There is a one-dimensional center manifold tangent to the right hand side of the ODE, F . This corresponds to $\lambda_2 \approx 0$, i.e., since clearly $F(x_1) \approx d\varphi_x F(x)$, the tangent vector roughly parallel to $F(x) \in T_x \mathbb{R}^3$ does not show exponential growth or decay under the tangent dynamics. Thus, this map is not uniformly hyperbolic as per the description in section 6.2.2. Rather, it is a partially hyperbolic system – a generalization of a uniformly hyperbolic system that allows a center direction – in which the center-unstable manifold is two-dimensional and tangent to $\text{span}\{F\} \oplus E^u$. The Lorenz attractor nevertheless mimics the statistical behavior of a uniformly hyperbolic attractor. For instance, the central limit theorem holds for Hölder continuous observables and an SRB-type invariant distribution exists [13].

In Figure 6-4, we numerically calculate the one-dimensional unstable manifold at $x := (0, 0, 1)$ of the Lorenz attractor. We populate the small line segment connecting $[-0.01, 0, 1]$ and $[0.01, 0, 1]$ with 10001 equi-spaced initial conditions. In Figure 6-4, these points are shown after time evolution for time $T_1 = 18$ or $n_1 = 1800$ steps (on the left) and $T_2 = 20$ or $n_2 = 2000$ steps (on the right). The points that are a small distance from one another at all times up to the indicated times are considered orbits within local unstable manifolds of the reference orbit $\{x_n\}$.

Along these selected orbits, we use the following finite difference approximation to compute V^1 :

$$V^1(y_n) \approx \frac{x_n - y_n}{\|x_n - y_n\|}. \quad (6.23)$$

The 3 components $V_{x_i}^1$, $i = 1, 2, 3$ obtained this way are shown in gray in Figure 6-5; to avoid confusing these scalar fields with $V^1(x_n)$, we do not use the shorthand notation, in this section, for $V^1(x_n)$, which refers to the first CLV at the phase point x_n . The scalar fields $V_{x_i}^1$ match match closely the results, shown in orange, of a more typical method of computing the first CLVs. This second method to compute $V^1(x_n)$ uses only the trajectory x, x_1, \dots, x_n and the tangent dynamics along this trajectory, and works as follows: randomly initialize $v(x)$ and propagate the tangent dynamics with

repeated normalization.

$$v(x_{n+1}) = d\varphi(x_n)v(x_n), \quad (6.24)$$

$$v(x_{n+1}) \longleftarrow v(x_{n+1})/\|v(x_{n+1})\|. \quad (6.25)$$

Carrying this out for $n \in \mathbb{Z}^+$, similar to a power iteration method for the computation of the dominant eigenvector of a matrix, yields a unit vector $v(x_n)$ that aligns with $V^1(x_n)$. As confirmed in Figure 6-5, this procedure is equivalent to the above-mentioned finite difference procedure, as long as y_n is in a small neighborhood of x_n , for the length of the trajectory considered.

Having visualized V^1 along trajectories, we now compute W^1 using our differential CLV method in section 6.3. To test its correctness, we also compute W^1 using a finite difference method as follows. As usual, let the reference trajectory along which we require to compute W^1 be x, x_1, \dots, x_N , and assume that we know the CLVs $V^1(x), V^1(x_1), \dots, V^1(x_N)$. Let y, y_1, \dots, y_N and r, r_1, \dots, r_N be two other trajectories that are at most a distance of $\mathcal{O}(1)$ away from the reference trajectory, at each of the N time steps. Then, according to our preceding discussion,

$$V^1(y_n) \approx -V^1(x_n) \approx \frac{x_n - y_n}{\|x_n - y_n\|}. \quad (6.26)$$

At each n , we rescale y_n and r_n along $V^1(x_n)$ to obtain the two points i) $\tilde{y}_n = x_n + \epsilon_{y_n} V^1(y_n)$, ii) $\tilde{r}_n = x_n + \epsilon_{r_n} V^1(r_n)$. Then, we can approximately compute $W^1(x_n)$ as

$$W^1(x_n) \approx \frac{(\tilde{r}_n - x_n)/\epsilon_{r_n} - (\tilde{y}_n - x_n)/\epsilon_{y_n}}{\|\tilde{r}_n - \tilde{y}_n\|}. \quad (6.27)$$

In Figure 6-6, we plot the three components of W^1 : $W_{x_1}^1, W_{x_2}^1, W_{x_3}^1$ computed using the above procedure in gray and the same quantity computed using the differential CLV algorithm in section 6.3 in orange. The closeness of the two results indicates the correctness of our algorithm. It is also a numerical verification of the fact that V^1 is differentiable along itself in this system, even though it is only partially hyperbolic.

6.4.3 Qualitative verification on a perturbed cat map

We consider a smoothly perturbed Cat map (PCM) (see section 6.2.3) due to Sli-pantschuk *et al.* [245]. The PCM [245] was designed to be an analytic, area-preserving, uniformly hyperbolic map of the torus, whose spectral properties can be computed analytically. The PCM is given by

$$\varphi([x_1, x_2]) = \begin{bmatrix} 2 & 1 \\ 1 & 1 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} + \begin{bmatrix} \Psi_{s_1, s_2}(x_1) \\ \Psi_{s_1, s_2}(x_1) \end{bmatrix}, \quad (6.28)$$

where

$$\Psi_{s_1, s_2}(y) := (1/\pi) \arctan \left(s_1 \sin(2\pi y - s_2) / (1 - s_1 \cos(2\pi y - s_2)) \right)$$

is a perturbation whose maximum magnitude is controlled by the parameter s_1 and the location of the maximum, by s_2 . Clearly, the original Cat map is recovered at $s_1 = 0$. As in the Cat map, the sum of the LEs is 0 but their values are sensitive to the parameters, with lesser sensitivity to s_2 when compared to s_1 . Unlike the Cat map, the CLVs are no longer uniform in phase space and are also not orthogonal to each other. In Figure 6-7, we show the vector fields V^1 and V^2 computed at $s_1 = 0.75$ and $s_2 = 0.2$. Notably, non-zero values of s_1 create a curvature in the CLVs, which is again non-uniform in space. We compute the self-derivative of the unstable CLV using our differential CLV method in section 6.3. By construction, the method produces a vector field W^1 that is orthogonal to V^1 . The norm of the computed vectors, $\|W^1\|$, is shown signed according to its orientation with respect to V^1 . In particular, in Figure 6-8, we plot $\|W^1 \times V^1\|$ as a colormap on the vector field V^1 . Figure 6-8 is a qualitative representation of the fact that $\|W^1\|$ is the curvature of the unstable manifold, which is everywhere tangent to the plotted vector field V^1 . The V^1 self-derivative W^1 is the acceleration of a particle moving with the velocity field V^1 . This intuitive picture is mirrored by Figure 6-8, in which $\|W^1\|$ is higher in regions of velocity changes than where the velocity appears rather uniform (e.g. in a thin strip around the diagonal of the square). The regions of similar magnitude of




Figure 6-7: The vector fields V^1 (left) and V^2 (right) are shown for the PCM at $s_1 = 0.75, s_2 = 0.2$.




Figure 6-8: The vector field V^1 is shown for the PCM at $s_1 = 0.75, s_2 = 0.2$. The color represents the values of $\|W^1 \times V^1\|$, which equals the norm of $\|W^1\|$ multiplied by a sign representing the orientation with respect to V^1 .

acceleration but of opposite sign, reflect the symmetry in the velocity field V^1 about $x_1 = x_2$, and moreover indicate the opposite directions of the *turns* made in those regions by traveling particles.

6.4.4 Qualitative verification on the volume-decreasing perturbed Cat

While the PCM was an example of a symplectic uniformly hyperbolic system, now we consider a dissipative uniformly hyperbolic map. We introduce another perturbed Cat map, with smooth nonlinear perturbations that cause the resulting map to be volume-decreasing. The norm of the perturbations is controlled by a set of four parameters $s = [s_0, s_1, s_2, s_3]^T$ and the unperturbed Cat map (the original Anosov Cat) is recovered at $s = [0, 0, 0, 0]$. The map, referred to as the dissipative Cat map or DCM hereafter, is defined as follows:

$$\begin{aligned} \varphi([x_1, x_2]^T) = & \begin{bmatrix} 2 & 1 \\ 1 & 1 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} + \left(s_0 \begin{bmatrix} v_0 \\ v_1 \end{bmatrix} + s_1 \begin{bmatrix} v_2 \\ v_3 \end{bmatrix} \right) \sin(2\pi \tilde{V}^2 \cdot x) / c \\ & + \left(s_2 \begin{bmatrix} v_0 \\ v_1 \end{bmatrix} + s_3 \begin{bmatrix} v_2 \\ v_3 \end{bmatrix} \right) \sin(2\pi \tilde{V}^1 \cdot x) / c \end{aligned} \quad (6.29)$$

where $\tilde{V}^2 := [v_0, v_1]^T = [5, -8]^T \in \mathbb{R}^2$ is a rational approximation of the stable CLV of the unperturbed Cat map. Similarly, $\tilde{V}^1 := [v_2, v_3]^T = [8, 5]^T \in \mathbb{R}^2$ is a rational approximation of the unstable CLV of the unperturbed Cat map. The constant c serves to normalize the perturbations and is set to $c = 2\pi(v_0^2 + v_1^2)$. The four parameters together determine the norm and direction of the perturbation. In Figure 6-9, V^1 is plotted in each case of turning on just one of the four parameters, in order to isolate its effects. Each subfigure reflects the effect of a single parameter on V^1 , in comparison to the unperturbed Cat map (in which V^1 is roughly parallel to the line \tilde{V}^1). For instance, when $s = [1, 0, 0, 0]^T$, a perturbation is applied along the direction \tilde{V}^2 , which is approximately along the stable direction of the DCM. The norm of this perturbation varies sinusoidally with the orientation along the approximately




Figure 6-9: The vector field V^1 is shown for the DCM at different parameter choices. The parameters not indicated are set to 0 in each case.

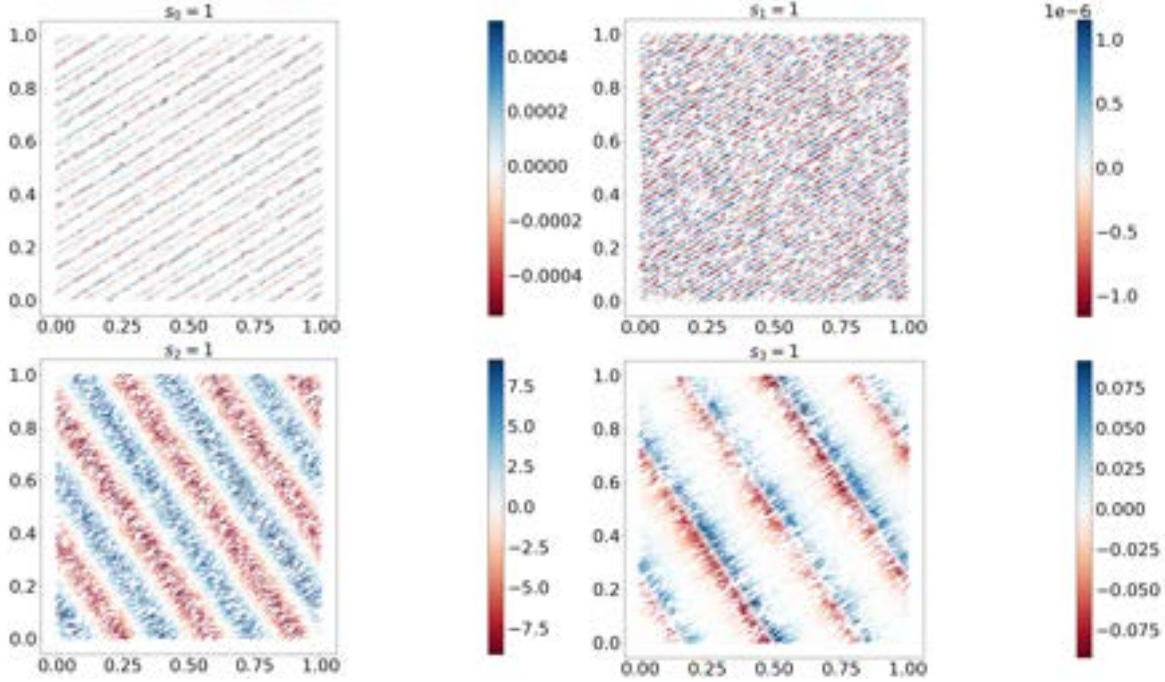


Figure 6-10: The vector field V^1 is shown for the DCM, colored according to $\|W^1 \times V^1\|$. The parameters not indicated as 1 are set to zero in each case.

stable direction, \tilde{V}^2 . As can be seen in the top-left of Figure 6-9, the CLV V^1 is rather uniform in its own direction but shows a striated pattern in the perpendicular direction, roughly along \tilde{V}^2 . As another example, the bottom-left subfigure shows V^1 at $s = [0, 0, 1, 0]^T$. From Eq. 6.29, we know that s_2 being non-zero introduces a perturbation, along \tilde{V}^2 , whose norm varies in the approximately unstable direction, \tilde{V}^1 . This is portrayed in the figure, wherein V^1 appears as waves, which are seen traveling approximately along \tilde{V}^2 but the amplitudes of the waves clearly vary in the perpendicular, approximately unstable direction. Turning on the parameter s_1 exchanges the roles of \tilde{V}^1 and \tilde{V}^2 when compared to when s_2 is non-zero. From the top-right subfigure in which the effect of s_1 is shown, we can see that there is no noticeable curving of the unstable manifold since the perturbation is aligned with the unstable direction. Finally, the effect of a non-zero s_3 is depicted in the bottom-right of Figure 6-9. Here we see the compression and expansion of unstable manifolds in the unstable direction since a perturbation non-uniform in the unstable direction is applied along the unstable direction.

With this understanding of the effect of each parameter, we expect that V^1 would show a smaller sensitivity, in its own direction, when the norm of the perturbation is uniform along \tilde{V}^1 . This is the case when s_2, s_3 are set to 0. This intuition is confirmed by the numerical results obtained on using the differential CLV method. As shown in Figure 6-10, when either $s_0 = 1$ or $s_1 = 1$, and the other 3 parameters are set to 0, we see that the numerically computed W^1 has a smaller norm, when compared to the other cases.

On the bottom row in Figure 6-10 are the vector fields W^1 when either s_2 or s_3 are set to 1 and the rest to 0. In these cases, the norm of the perturbation varies along the approximately unstable direction, and this is clearly reflected in the higher (when compared to the other two cases) magnitudes of W^1 . In addition, the variation in W^1 itself, which gives information about the second-order derivative of V^1 , is also consistent with our expectations. For instance, W^1 shows a marked variation along V^1 when $s_2 = 1$ (bottom-left of Figure 6-10). This can be explained by the applied perturbations being sinusoidal in the direction of \tilde{V}^1 , giving rise to a harmonic functions for the higher-order derivatives along V^1 as well. Finally, when $s_3 = 1$, (bottom-right of Figure 6-10), it is easy to observe that, qualitatively, the *density* of the lines V^1 is reflected in the magnitudes of W^1 . This is not a coincidence, as we shall see in section 6.5. There, we describe that W^1 is indirectly related to the variation in the density of the SRB measure on the unstable manifold, due to perturbations along V^1 . Now we can see that especially the $s_3 = 1$ case provides a visualization consistent with this theoretical insight. Particularly, the pronounced variation in the unstable direction (bottom-right, Figure 6-10), mirrors the changes in probability density on the unstable manifold, which is qualitatively measured by the closeness of the V^1 lines in Figure 6-9.




Figure 6-11: The CLV V^1 on the henon attractor. Inset is the CLV field in a neighborhood of the fixed point $\approx (0.63, 0.19)$.

6.4.5 Numerical results on the Hénon map

As our final example, we consider the classical Hénon attractor. The Hénon map is the canonical form for a two-dimensional area-decreasing quadratic map [135]:

$$\varphi([x_1, x_2]^T) = \begin{bmatrix} x_2 + 1 - s_0 x_1^2 \\ s_1 x_1 \end{bmatrix}. \quad (6.30)$$

Taking the parameters s_0 and s_1 at their standard values of $s_0 = 1.4$ and $s_1 = 0.3$, we obtain the Hénon attractor, on which the CLVs are shown in Figure 6-11. At these parameter values, the Hénon attractor is nonhyperbolic due to the presence of tangencies between the stable and unstable manifolds [12]. On this map, we apply the differential CLV method we derived in section 6.3, and the resulting W^1 is shown in Figure 6-12. The CLVs may not be differentiable everywhere, as seen by the large magnitudes of the numerically computed W^1 at the sharp turns in the attractor. In Figure 6-13, we dissect the derivatives further to investigate the issue of differentiability numerically. In each subfigure, the vector field V^1 is plotted colored according to $\|W^1\|$; at the points at which $\|W^1\|$ is not in the range indicated by the colormap, V^1 is shown using thin black lines. From the top row of Figure 6-13, it is clear that $\|W^1\| < 0.1$ for the relatively straight portions of the attractor and the points on the right, curved side of the attractor, still have a curvature less than 1. On the bottom row, the more rounded portions of the attractor, as expected, have a higher curvature




Figure 6-12: The vector field V^1 is shown for the Hénon map. The color represents the V^1 self-derivative norm, $\|W^1\|$.




Figure 6-13: The vector field V^1 is shown for the Hénon map. The color represents $\|W^1\|$, the curvature of the unstable manifold.

when compared to the previous cases. On the bottom-right, we see that only the corners and turns have $\|W^1\|$ higher than 100. Among these points, the variation in the curvature, $\|W^1\|$, is over six orders of magnitude, with the sharp corners having the highest curvatures. In this case, our numerical method for W^1 acts as an indicator for the lack of differentiability at some points. At least in two dimensions, this also turns out to be a detector for uniform hyperbolicity, based on our discussion in section 6.2.4.

6.5 An application of CLV derivatives to statistical linear response

A landmark result in the theory of uniformly hyperbolic systems due to Ruelle ([232][234] ; [118] contains a modern proof of the result) is the smooth response of their statistics to parameter perturbations. Here we briefly describe this result, called the linear response formula, and draw a connection between the formula and Eq. 6.17, which is the differential expansion equation.

Consider a family of uniformly hyperbolic maps $\varphi_s \in C^3(\mathbb{M})$, where s is a small parameter around 0. Let the reference map φ_0 be written simply as φ , and V be a smooth vector field such that $\varphi_s = \varphi + sV$ up to first order in s . Let the SRB measure of φ_s be μ_s : that is, μ_s is a φ_s -invariant probability distribution on \mathbb{M} such that for any continuous scalar observable J , the ergodic average starting from a $x \in \mathbb{M}$ Lebesgue-a.e., $\lim_{N \rightarrow \infty} (1/N) \sum_{n=0}^{N-1} J(\varphi_s^n(x)) = \langle J, \mu_s \rangle$.

Ruelle's linear response theory [232][234] proves the existence of the statistical response to parameter changes, $\langle J, \partial_s \mu_s \rangle$, in uniformly hyperbolic systems, including expressing this quantity as an exponentially converging series, which is known as linear response formula. The quantity $\langle J, \partial_s \mu_s \rangle$ represents the derivative with respect to s of ergodic averages or equivalently ensemble averages of observables with respect to the SRB measure, and is of immense interest in practical applications. The statistical sensitivity $\langle J, \partial_s \mu_s \rangle$ is useful for sensitivity analysis, uncertainty quantification,

model selection etc, in every scientific discipline from climate studies [228][185] to aerodynamic fluid flows [208][142][71]. The linear response formula [232][234] is as follows:

$$\langle J, (\partial_s \mu_s) \Big|_0 \rangle = \sum_{n=0}^{\infty} \langle d(J \circ \varphi^n) \cdot V, \mu_0 \rangle. \quad (6.31)$$

Although the above series is exponentially converging, previous works [69][96] suggest that it is computationally infeasible to calculate the series in its original form when V has a non-zero component in E^u , especially in high-dimensional practical systems. This is because the integrand in each term increases exponentially with n : $|d(J \circ \varphi^n) \cdot V| \sim \mathcal{O}(\exp(\lambda_1 n))$, for almost every perturbation V , which will have a non-zero component along V^1 . If each term in the series is regularized by an integration by parts, the resulting form of the linear response formula is more amenable to computation.

For a simple illustration, we consider the case of one-dimensional unstable manifolds, and fix the smooth perturbation field to be $V = a V^1$, which has a scalar component, a , along the unstable CLV. Applying integration by parts to Eq. 6.31 on the unstable manifold [232][234] (see also Appendix section 6.8.1), and then using the fact that ergodic averages converge to ensemble averages for Lebesgue-a.e. x ,

$$\langle J, (\partial_s \mu_s) \Big|_0 \rangle = - \sum_{k=0}^{\infty} \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{n=0}^{N-1} J(x_{k+n}) (a(x_n) g(x_n) + b(x_n)) \quad (6.32)$$

where

- ρ_0 is the density of the conditional distribution of μ_0 on unstable manifolds [281];
- $g(x) := \frac{1}{\rho_0(x)} \frac{d(\rho_0 \circ \mathcal{C}_{x,1})(t)}{dt} \Big|_{t=0}$, is the *logarithmic density gradient* function; and,
- $b(x) := \frac{d(a \circ \mathcal{C}_{x,1})(t)}{dt} \Big|_{t=0}$, is the derivative of a along unstable manifolds.

The computational infeasibility of Ruelle's original expression in Eq. 6.31 is overcome

by Eq. 6.32, as it results from regularization through integration by parts. That is, the ergodic averaging computation, listed in Eq. 6.32, follows the central limit theorem, with an error convergence as $\mathcal{O}(1/\sqrt{N})$, when computed along an orbit of length N . To compute Eq. 6.32, we must determine the two functions g and b along orbits. The derivative b can be computed at any x_n as $b(x_n) = (V_n^1)^T(dV)_n V_n^1$; to derive this expression, we use the fact that $W^1 \cdot V^1 = 0$, which follows from Eq. 6.19. Now, the only other unknown is the fundamental quantity – the logarithmic density gradient denoted g . Using the fact that φ preserves μ_0 , it can be shown that (see section 4 of [62] for an alternative derivation and [247] for an intuitive description of g on one-dimensional unstable manifolds) g satisfies the following iterative equation along trajectories:

$$g_{n+1} = \frac{g_n}{z_{n,1}} + \alpha_{n,1}. \quad (6.33)$$

In the above equation, we use the shorthand notation $g_n := g(x_n)$ and $z_{n,1} := z_{x_n,1}$, fixing any μ_0 -typical x . Thus, Eq. 6.33 is an iterative formula that can be used to compute g along orbits. It uses differential expansion equation (Eq. 6.17) for the second term on the right hand side. The values of g along a typical orbit, thus computed, are used in Eq. 6.32 to obtain the desired sensitivity.

6.6 Conclusion

In this work, we have derived a numerical method, called the differential CLV method, to compute the derivatives of Covariant Lyapunov Vectors along their own directions: the CLV self-derivatives. These directional derivatives exist in smooth uniformly hyperbolic systems with compact attractors. The differential CLV method converges asymptotically at an exponential rate in the case of the CLV self-derivatives corresponding to the largest and smallest Lyapunov exponents. We demonstrate the application of the differential CLV method on a variety of systems with one-dimensional unstable manifolds including a quasi-hyperbolic attractor (Lorenz'63) and a non-

hyperbolic attractor (Hénon). In the two-dimensional uniformly hyperbolic systems considered, including perturbations of the Cat map, our method provides rich visualizations of the curvature of the one-dimensional unstable manifold. A byproduct of the differential CLV method, without the orthogonal projection step (Eq. 6.19), known as the differential expansion equation (Eq. 6.17), is fundamentally linked to the statistical linear response of a chaotic attractor. The link is through its utility to compute the divergence of perturbations on the unstable manifold, with respect to the SRB measure conditioned on unstable manifolds. This connection makes the differential expansion derivatives concretely useful for efficiently differentiating statistics with respect to system parameters in uniformly hyperbolic systems. It must be explored in a future work how the differential CLV method can be made to have unconditional asymptotic convergence for the self-derivatives of all CLVs, as opposed to only the most unstable and the most stable CLVs, which are treated in this work. With this generalization, the second-order tangent equations presented in this paper can spawn applications to sensitivity analysis in chaotic systems, and beyond.

6.7 Appendix

6.7.1 The lack of differentiability of CLVs

In general, we say that a subspace E is Hölder continuous on \mathbb{M} if there exist constants $K, \delta > 0$ and $\beta \in (0, 1]$ such that $\|E_x - E_y\|_* \leq K \|x - y\|^\beta$, whenever $x, y \in \mathbb{M}$ are such that $\|x - y\| \leq \beta$. As mentioned in section 6.2.4, the subspaces E^u, E^s are Hölder continuous spaces with an β that is rarely equal to 1. The reader is referred to classical texts such as [154] (Chapter 19) or [130] for a detailed exposition on Hölder structures on hyperbolic sets.

There, the norm $\|\cdot\|_*$ uses an adapted coordinate system such as the one introduced in section 6.3.1. The set of Hölder continuous functions themselves, is independent of the coordinate system, however. The norm $\|\cdot\|_*$ used in the above references (e.g. in Theorem 19.1.6 of [154]), for our particular choice of adapted co-

ordinates introduced in 6.3.1, results in the following definitions, which are exactly what one might expect. Suppose $\|x - y\| \leq \delta$, and Q_x, Q_y are matrix representations of the CLV basis whose i th columns respectively are V_x^i, V_y^i . Then, $\|E_x^u - E_y^u\|_* := \|Q_x[:, 1 : d_u] - Q_y[:, 1 : d_u]\|$ where the norm on the right hand side is a matrix norm on $\mathbb{R}^{m \times d_u}$, say the induced 2-norm. Here we have again used programmatic notation: given a matrix A , $A[:, i : j]$ refers to the columns of A from i to j , limits included. Similarly, for E^s , $\|E_x^s - E_y^s\|_* := \|Q_x[:, d_u + 1 : d] - Q_y[:, d_u + 1 : d]\|$. Consistent with these definitions, for a one-dimensional E^i , we have $\|E_x^i - E_y^i\| := \|V_x^i - V_y^i\|$, which is simply the 2-norm on \mathbb{R}^m .

6.7.2 Computations on the super-contracting Solenoid attractor

The super-contracting Solenoid attractor is the curve $\gamma : [0, 2\pi] \rightarrow \mathbb{R}^3$ (defined in Eq. 6.21) parameterized by a single parameter t . Since we have a closed form expression for the one-dimensional attractor, we can compute its tangent vector field, as:

$$\frac{d\gamma}{dt} = \begin{bmatrix} -2r_1(t) \sin 2t - (\sin t \cos 2t)/2 \\ 2r_1(t) \cos 2t - (\sin t \sin 2t)/2 \\ \frac{\cos t}{2} \end{bmatrix}, \quad (6.34)$$

where

$$r_1(t) = \left(s_0 + \frac{\cos t}{2} \right).$$

As explained in section 6.4.1, $V^1(t) = \gamma'(t) / \|\gamma'(t)\|$. Further, we analytically calculate that

$$\partial_{\gamma'(t)/\|\gamma'(t)\|} (\gamma'(t)/\|\gamma'(t)\|) = \frac{1}{2} \begin{bmatrix} -(193 \cos t + 392 \cos 2t + 267 \cos 3t + 68 \cos 4t + 6 \cos 5t + 36)/c_1 \\ -(189 \sin t + 392 \sin 2t + 267 \sin 3t + 68 \sin 4t + 6 \sin 5t)/c_1 \\ -(19 \sin t + 8 \sin t \cos t + 2 \sin t \cos 2t - 2 \sin 2t \cos t)/(c_1/2) \end{bmatrix} \begin{bmatrix} -\sin 2t/r_1 & \cos 2t/r_1 & 0 \end{bmatrix} \frac{\gamma'(t)}{\|\gamma'(t)\|} \quad (6.35)$$

where

$$c_1 := 2(16 \cos t + 2 \cos 2t + 19)^{3/2}.$$

In Figures 6-3 and 6-2, we observe that the vector field W^1 computed using the differential CLV method (Eq. 6.19), matches almost exactly against the above expression in Eq. 6.35.

6.8 Convergence of the differential CLV method

In this section, we show that convergence of Eq. 6.19 is guaranteed when $i = 1$. Moreover, the asymptotic convergence is exponentially fast. Fix a reference trajectory q, q_1, \dots , and use the notation f_n to denote $f(x_n)$. Let W^i, W_1^i, \dots and $\tilde{W}^i, \tilde{W}_1^i, \dots$ be two sequences of vectors generated by iterating Eq. 6.19. Then, from Eq. 6.19,

$$\left\| W_n^i - \tilde{W}_n^i \right\| = \frac{1}{\prod_{m=0}^{n-1} z_{m,i}^2} \left\| \prod_{m=0}^{n-1} \left((I - V_{m+1}^i (V_{m+1}^i)^T) (d\varphi)_m \right) (W^i - \tilde{W}^i) \right\|. \quad (6.36)$$

We can apply Oseledets MET to the cocycle $\text{Coc}(x_m, n) = \prod_{k=0}^{n-1} (I - V_{m+k+1}^i (V_{m+k+1}^i)^T) (d\varphi)_{m+k}$, and to the Jacobian cocycle to obtain the following asymptotic inequality. In particular, using the relationship Eq. 6.6, we get that for every $\epsilon > 0$, there exists an $N \in \mathbb{N}$ such that for all $n \geq N$,

$$\begin{aligned} \left\| W_n^i - \tilde{W}_n^i \right\| &= \frac{1}{\prod_{m=0}^{n-1} (z_m^i)^2} \left\| \prod_{m=0}^{n-1} \left((I - V_{m+1}^i (V_{m+1}^i)^T) (D\varphi)_m \right) (W^i - \tilde{W}^i) \right\| \\ &\leq e^{-2n(\lambda_i - \epsilon)} e^{n(\omega_i + \epsilon)} \left\| W^i - \tilde{W}^i \right\|. \end{aligned} \quad (6.37)$$

In the above inequality 6.37, $\omega_i := \max_{j \neq i, 1 \leq j \leq d_u} \lambda_j$. Thus, asymptotic exponential convergence is guaranteed whenever $2\lambda_i \geq \omega_i$, which is of course true when $i = 1$.

6.8.1 Regularization of Ruelle's formula

Here we briefly describe the derivation of Eq. 6.32 from Ruelle's formula (Eq. 6.31). The reader is referred to Ruelle's original papers [232][234], or to [62] for an alternative

derivation of a regularized response to unstable perturbations. In the case of one-dimensional unstable manifolds, which is the focus of this paper, we can obtain Eq. 6.32 by the following sequence of steps:

- Disintegration of the SRB measure on the unstable manifolds. Let Ξ be a partition of \mathbb{M} subordinate to the unstable manifold [169], and let ρ_0 be the conditional density of the SRB measure on elements of Ξ . Then, disintegration results in the following expression for the (n th term in the) linear response to the unstable perturbation aV^1 ,

$$\langle d(J \circ \varphi^n) \cdot a V^1, \mu_0 \rangle = \int_{\mathbb{M}/\Xi} \int_{\Xi(x)} \left(d(J \circ \varphi^n) \cdot a V^1 \right) \circ \mathcal{C}_{x,1}(t) \rho_0 \circ \mathcal{C}_{x,1}(t) dt d\hat{\mu}_0(x) \quad (6.38)$$

$$= \int_{\mathbb{M}/\Xi} \int_{\Xi(x)} a \circ \mathcal{C}_{x,1}(t) \frac{d(J \circ \varphi^n \circ \mathcal{C}_{x,1})}{dt}(t) \rho_0 \circ \mathcal{C}_{x,1}(t) dt d\hat{\mu}_0(x). \quad (6.39)$$

In the above expression, $\Xi(x)$ is the element of Ξ containing x , and the quotient measure of the SRB measure on M/Ξ is denoted $\hat{\mu}_0$.

- Applying integration by parts on the inner integral, we obtain,

$$\langle d(J \circ \varphi^n) \cdot a V^1, \mu_0 \rangle = \int_{\mathbb{M}/\Xi} \int_{\Xi(x)} \frac{d((a \rho_0 J \circ \varphi^n) \circ \mathcal{C}_{x,1})}{dt}(t) dt d\hat{\mu}_0(x) \quad (6.40)$$

$$- \int_{\mathbb{M}/\Xi} \int_{\Xi(x)} J \circ \varphi^n \circ \mathcal{C}_{x,1}(t) \left(\frac{a \circ \mathcal{C}_{x,1}(t)}{\rho_0 \circ \mathcal{C}_{x,1}(t)} \frac{d(\rho_0 \circ \mathcal{C}_{x,1})}{dt}(t) \right) dt d\hat{\mu}_0(x) \quad (6.41)$$

$$+ \frac{d(a \circ \mathcal{C}_{x,1})}{dt}(t) \rho_0 \circ \mathcal{C}_{x,1}(t) dt d\hat{\mu}_0(x). \quad (6.42)$$

The first term on the right hand side of the above equation vanishes, as noted by Ruelle [232][234] for arbitrary dimensional unstable manifolds in Theorem 3.1(b). Applying the divergence theorem on the first term, we obtain integrals over boundaries of the partition elements, which incur cancellations in the outer integral.

- Using the definitions of b and g in the above equation, we obtain

$$\langle d(J \circ \varphi^n) \cdot a V^1, \mu_0 \rangle = -\langle J \circ \varphi^n (a g + b), \mu_0 \rangle. \quad (6.43)$$

Eq. 6.32 is now obtained when we rewrite the above ensemble average as an ergodic average.

Chapter 7

Space-split sensitivity algorithm for one-dimensional unstable manifolds

Our earlier attempt at the computation of the Ruelle's formula (Chapter 5) and our numerical procedure to differentiate the first CLV in the previous chapter culminate in the space-split sensitivity or the S3 algorithm in this chapter. The S3 algorithm transforms Ruelle's formula into a well-conditioned ergodic-averaging computation. *The decomposition of Ruelle's formula, and its subsequent transformation, are different from that obtained by stable-unstable splitting of the perturbation field, which is used in Chapter 5.* The main motivation for deriving a new decomposition of Ruelle's formula, henceforth called the S3 decomposition, is its differentiability on the unstable manifold, which is lacking for the stable-unstable splitting of the perturbation field. This differentiability on the unstable manifold of the two resulting vector fields that arise from the S3 decomposition is key to deriving an efficient, trajectory-based computation.

The S3 decomposition is based on the idea of repeatedly orthogonalizing the conventional tangent equation. One of the resulting terms, which in an abuse of terminology is again called the stable contribution, can be computed using a regularized tangent equation, similar to in a non-chaotic system. The remaining term, known as the unstable contribution, is regularized and converted into an efficiently computable ergodic average.

In this process, we develop new algorithms, which may be useful beyond linear response, to compute the unstable derivatives of the regularized tangent vector field and the unstable direction. We prove that the S3 algorithm, which combines these computational ingredients that enter the stable and unstable contributions, converges like a Monte Carlo approximation of Ruelle’s formula. The algorithm presented here is hence a first step toward full-fledged applications of sensitivity analysis in chaotic systems, wherever such applications have been limited due to lack of availability of long-term sensitivities. The background (section 7.2), targeted at computational scientists from across disciplines, aims at concisely providing the mathematical preliminaries concerning dynamical systems and ergodic theory used in this thesis. This chapter is part of a submitted work [67].

7.1 Introduction

We say that a parameterized family of dynamical systems obeys *linear response* when the infinite-time averages or ergodic averages of its smooth observables vary differentiably with the parameter. It was shown by Ruelle [232][234] that uniformly hyperbolic maps, which are mathematical idealizations of chaotic attractors, follow linear response; Proposition 8.1 of [118] is a simplified proof using modern transfer operator techniques. Rigorous proofs of linear response have since been extended to uniformly hyperbolic flows [235], partially hyperbolic systems [85], dissipative stochastic systems [126], stochastically perturbed uniformly hyperbolic systems [181] and even to a larger class of stochastic systems with possibly non-hyperbolic unperturbed dynamics [106], certain nonuniformly hyperbolic systems [44] and intermittent systems [29][25][21]. From the statistical physics point of view, linear response theory has been found to be robust in high-dimensional systems [277][276], and has been usefully applied to chaotic systems across disciplines including climate models [228][223][50], biological systems (see [55] for a review of linear response in neuronal networks) and turbulent flows in engineering systems [71][37][208][257].

Linear response of a chaotic system quantifies the proportional change in its long-

term statistical behavior in response to small parameter perturbations. Apart from providing phenomenological understanding, this measure of long-term sensitivity is immensely useful in practical chaotic systems, which are often high-fidelity numerical simulations, for computational applications such as optimization, uncertainty quantification and parameter selection [83][37][41][71]. Particularly in climate science, theoretical as well as computational studies of violations of linear response and the presence of arbitrarily large linear responses [183][186][73][236] are crucial to gain a better understanding of intermittencies and climate tipping points [18], which are active areas of research.

Ruelle [232][234] established a formula for linear response or the derivative of ergodic averages with respect to parameters. However, a direct evaluation of this formula typically shows a poor convergence rate and is computationally impractical, as previous works have shown [69][96]. This poor convergence is due to the exponential growth of infinitesimal perturbations – the so-called butterfly effect – which is the defining characteristic of chaotic systems. Due to the butterfly effect, the sensitivity of a state at a time n into the future, to an infinitesimal perturbation to the current state in almost any direction, grows exponentially with n . Now, infinitesimal parameter perturbations may be thought of infinitesimal perturbations to each state applied in different tangent directions. Thus, along any trajectory, the sensitivity to an infinitesimal parameter perturbation also grows exponentially. However, the average of the sensitivities across all trajectories is a bounded quantity at all times. This ensemble-averaged sensitivity is, in fact, exponentially decreasing with time in uniformly hyperbolic systems, and Ruelle’s formula for linear response is a series summation of these ensemble-averaged sensitivities. But, ensemble averaging exponentially growing quantities is a computationally challenging task that shows poor convergence.

Traditionally, sensitivities in a dynamical system are estimated by using tangent or adjoint equation solutions, or through automatic differentiation. Since all these methods time-evolve infinitesimal perturbations about a reference trajectory, the sensitivities they compute grow exponentially along any trajectory. Hence, conventional

methods for sensitivity analysis have long been recognized as unsuitable in chaotic systems for computing linear response. Some methods circumvent the problem of exponentially growing sensitivities to compute a bounded value for linear response, but they may exhibit a bias [166][206]. One such method is least-squares shadowing [271][211], in which the shadowing lemma (see e.g. Chapter 18 of [154]) is used to compute sensitivities along a shadowing orbit. However, since shadowing orbits may be nonphysical, i.e., ergodic averages along shadowing orbits may not converge to ensemble averages, the sensitivities computed by least-squares shadowing are not guaranteed to converge to linear response [61][209].

Moreover, the direct evaluation of Ruelle’s formula may also be thought of as an adaptation of conventional tangent/adjoint-based sensitivity computations for chaotic systems. This direct evaluation, known as the ensemble sensitivity approach [168][96], involves taking a sample average of sensitivities computed by conventional tangent/adjoint methods. However, as we noted earlier, the number of samples needed, to reduce the variance in the exponentially growing sensitivities and compute linear response accurately, makes this approach computationally infeasible. In blended response algorithms [1], the ensemble sensitivity approach for short-time sensitivities is *blended* with a fluctuation-dissipation theorem-based approximation of the long-term sensitivities. This approximation is however adhoc since the densities of the SRB measure on unstable manifolds may not follow the fluctuation-dissipation theorem-based approximation, even though linear response holds.

The purpose of this chapter is a rigorous computation of linear response in chaotic dynamical systems. We remark that using transfer operator techniques, a rigorous computation of linear response has been developed before, but it has been restricted to low-dimensional expanding maps [22]. Our aim is to develop a numerical method to evaluate Ruelle’s formula that is scalable to high-dimensional practical systems. For this reason, we seek a method to compute Ruelle’s formula that is provably convergent and is a computable ergodic average, which does not involve discretization of the phase space. The latter property of trajectory-based computation ensures that the convergence is that of a Monte Carlo computation of Ruelle’s formula, at a rate

independent of the system dimension.

In this chapter, we develop the space-split sensitivity or the S3 method, which is a scalable, efficient and rigorous computation of linear response. We prove that S3 provably converges in uniformly hyperbolic systems, and the convergence rate is similar to a typical Monte Carlo integration. We focus on uniformly hyperbolic systems with one-dimensional unstable manifolds, but design S3 keeping in mind future extensions to systems that have an unstable manifold of arbitrary dimension. To derive S3, we prove a decomposition of Ruelle’s formula, which is differentiable on the unstable manifold, into stable and unstable contributions to the overall sensitivity. Although we refer to the resulting components of linear response as stable and unstable contributions, note that this is not a stable-unstable splitting of the parameter perturbation vector field (Chapter 5), which is known to not be differentiable on the unstable manifold [221][130]. We develop recursive algorithms to compute the ingredients of both these components.

The plan for the chapter is as follows. In section 7.2, we introduce Ruelle’s linear response formula and provide the mathematical background for its decomposition and subsequent evaluation via the S3 algorithm. Section 7.3 is a concise statement of the main contributions of this chapter: two theorems concerning the S3 decomposition and evaluation, and the S3 algorithm. In section 7.4, we derive the S3 decomposition of Ruelle’s formula. The computation of the stable contribution that results from the decomposition is discussed in 7.4.3. An alternative expression of the unstable contribution is derived in section 7.4.4, whose computation is tackled in section 7.5. The validation of S3 on perturbed Baker’s maps is presented in section 7.6. The proofs of the two main theorems are split between sections 7.7 and 7.8. Section 7.7 proves the existence of the S3 decomposition and its differentiability in the unstable direction, while section 7.8 completes the proof of convergence of the S3 algorithm. In section 7.9, we summarize our contributions and present a roadmap for extending the present algorithm to systems with higher-dimensional unstable manifolds.

7.2 Preliminaries and problem setup

Consider a parameterized family $\varphi_s : M \rightarrow M$ of C^3 diffeomorphisms of a Riemannian manifold M , which we consider to be specified as a subset of \mathbb{R}^m . Let s be a scalar parameter that can take a small range of values around a reference value, say s_0 . Corresponding to infinitesimal parameter perturbations at s_0 , we define a vector field χ , whose value at $x \in M$ is given by $\chi_x := \partial_s \varphi_s(\varphi_s^{-1}x, s_0)$.

Suppose φ_s exhibits linear response at the reference value s_0 , which we associate with the unperturbed dynamics. This means that the parametric derivative of long-term averages of φ_s exists at the reference value s_0 . Thus, at any $s = s_0 + \delta s$ close to s_0 , up to first order in δs , the long-term averages of φ_s can be expressed using information associated only to the dynamics φ_{s_0} . Let $J \in C^2(M)$ be an observable of interest, e.g. a lift or a drag in a numerical simulation of a turbulent flow. We are interested in a quantitative determination of how the long-term average of J responds to infinitesimal perturbations in s , at s_0 , i.e., to perturbations in the direction χ .

7.2.1 Ergodic theory and linear response

The infinite-time average of J is defined as $\langle J \rangle(x) := \lim_{N \rightarrow \infty} (1/N) \sum_{n=0}^{N-1} (J \circ \varphi_s^n)(x, s)$, for $x \in M$. A probability distribution over the states on M , μ_s is stationary or φ_s -invariant, when $\mu_s(\varphi_s^{-1}A) = \mu_s(A)$, for any Borel subset $A \subseteq M$. When μ_s is an ergodic, φ_s -invariant, physical probability distribution for φ_s , infinite-time averages, also known as ergodic averages, are equal to expectations with respect to μ_s , at Lebesgue almost every x in the basin of attraction of φ_s . The expectation or ensemble average of an observable J with respect to μ_s , which is the Lebesgue integral with respect to the distribution μ_s , is written as $\langle J, \mu_s \rangle$.

That is, in ergodic systems, $\langle J \rangle(x, s)$ is independent of x – it is only a function of s – and equal to $\langle J, \mu_s \rangle$ at almost every x in a set of full Lebesgue measure. The quantity we wish to compute is *linear response* at s_0 , which is the parametric derivative of the ergodic/ensemble average, $d_s|_{s_0} \langle J, \mu_s \rangle$. When J explicitly depends on s and the dependence is smooth, $d_s \langle J, \mu_s \rangle = \langle J, \partial_s \mu_s \rangle + \langle \partial_s J, \mu_s \rangle$. The second

term, $\langle \partial_s J, \mu_s \rangle$ is simply an ergodic/ensemble average of $\partial_s J$, which can be computed, using its definition, along almost every orbit:

$$\langle \partial_s|_{s_0} J, \mu_{s_0} \rangle = \lim_{N \rightarrow \infty} (1/N) \sum_{n=0}^{N-1} (\partial_s J)(\varphi_{s_0}^n x, s_0),$$

for Lebesgue almost every x on M . Since this does not pose a computational challenge, we neglect the explicit dependence of J on s , and focus on the first term, $d_s|_{s_0} \langle J, \mu_s \rangle = \langle J, \partial_s|_{s_0} \mu_s \rangle$. From here on, we denote linear response at s_0 simply as $\langle J, \partial_s \mu_s \rangle$, for brevity.

7.2.2 Tangent dynamics

The tangent space at $x \in M$, denoted $T_x M$, is the space of all infinitesimal perturbations to x , which can be identified with \mathbb{R}^m . A vector field on M , which is a direction of infinitesimal perturbation to the state at each point on M , can be considered as a map from M to \mathbb{R}^m . If v is a vector field and $x \in M$, then $v_x \in T_x M \equiv \mathbb{R}^m$ denotes the value of the vector field at x . Fixing s at its reference value, we denote φ the map φ_{s_0} . Similarly, we refer to the distribution μ_{s_0} simply as μ .

The matrix $d\varphi^n$ gives the pushforward of a vector field by φ^n . That is, $w = d\varphi^n v$ if $w_{\varphi^n x} = d\varphi_x^n v_x$. We write $d\varphi^1$ simply as $d\varphi$. In the context of computations, we often fix a particular reference trajectory, say $\{x_n\}$, where each x_n is sampled according to μ . We denote values of a vector field v along the reference trajectory using $v_n := v_{x_n}$ for short; similarly, we write $(d\varphi)_n$ to denote the value of $d\varphi$ at x_n . The homogeneous tangent equation tracks the pushforward by φ^n along a fixed trajectory,

$$u_{n+1} = (d\varphi)_n u_n, \quad n \in \mathbb{Z}^+. \quad (7.1)$$

At every iteration of the homogeneous tangent equation, starting with $u_0 \neq 0 \in \mathbb{R}^m$, the vector field u is updated to $d\varphi u$. When we add a source term to Eq. 7.1, we refer to the resulting equation as the inhomogeneous tangent equation. For example, when the source term is the parameter perturbation field, $\chi := (\partial_s \varphi_s)(\varphi_{s_0}^{-1} \cdot, s_0)$, the

inhomogeneous tangent equation is the conventional tangent equation that is standard in sensitivity analysis,

$$u_{n+1} = (d\varphi)_n u_n + \chi_{n+1}. \quad (7.2)$$

In this tangent equation, at every iteration, the vector field u is updated to $d\varphi u + \chi$. In tangent sensitivity analysis, the parametric derivative of time-averages, $\partial_s(1/N) \sum_{n=0}^{N-1} (J \circ \varphi_s^n)(x_0, s_0)$, is usually computed using inhomogeneous tangent solutions (Eq. 7.2) as

$$\frac{1}{N} \sum_{n=0}^{N-1} \partial_s (J \circ \varphi_s^n)(x_0, s_0) = \frac{1}{N} \sum_{n=0}^{N-1} dJ(x_n) \cdot u_n.$$

Everywhere, d denotes the gradient operator in \mathbb{R}^m . For example, in the above equation, $dJ(x_n) = (dJ)_n \in \mathbb{T}_{x_n}^* M \equiv \mathbb{R}^m$ refers to the gradient of J evaluated at the point x_n .

7.2.3 Chaotic systems

The Oseledets multiplicative ergodic theorem (OMET) says that the homogeneous tangent solutions (Eq. 7.1) grow/decay asymptotically. Further, this asymptotic growth/decay is exponential at a finite number of rates, called the Lyapunov exponents, which, in our setting of ergodic systems, are independent of the starting point x_0 , at μ_s almost every x_0 . We denote them in descending order as $\lambda_1 > \lambda_2 > \dots > \lambda_p$, with $p \leq d$. A chaotic system, by definition, exhibits at least one positive Lyapunov exponent.

The tangent space at μ -almost every point $x \in M$ has a direct sum decomposition, $T_x M = \oplus_{i \leq p} E_x^i$, with the subspace E_x^i of tangent vectors having the asymptotic growth/decay rate λ_i . In other words, every tangent vector v_x belonging to the subspace E_x^i is such that $\lim_{n \rightarrow \infty} (1/n) \log \|d\varphi_x^n v_x\| = \lambda_i$. At μ -almost every x_0 , almost every choice of $u_0 \in \mathbb{R}^m$ will have a non-zero component on the tangent subspace corresponding to the largest positive Lyapunov exponent. Hence, in chaotic systems, homogeneous tangent solutions starting from every initial condition u_0 grow expo-

nentially in time. That is, for almost every $u_0 = u(x_0)$, $\|u_n\| \sim \mathcal{O}(e^{\lambda_1 n})$, with $\lambda_1 > 0$, for large n . The exponent λ_1 , which determines the asymptotic growth factor of the tangent solutions, is the largest among the Lyapunov exponents.

7.2.4 Uniform hyperbolicity

In this chapter, we consider our C^3 diffeomorphism $\varphi : M \rightarrow M$ to be equipped with a compact, invariant, hyperbolic attractor Λ , which contains a one-dimensional unstable manifold. That is, Λ is a compact set and $\varphi(\Lambda) = \Lambda$. We say Λ is a hyperbolic set for φ if there exist constants $C > 0$ and $\lambda \in (0, 1)$ such that, at every point $x \in M$, the tangent space $T_x M$ has a direct sum decomposition $T_x M = E_x^u \oplus E_x^s$, where

- E_x^u and E_x^s are φ -invariant, or *covariant*, subspaces. That is, $d\varphi_x(E_x^u) = E_{\varphi(x)}^u$ and $d\varphi_x(E_x^s) = E_{\varphi(x)}^s$.
- E_x^u is the 1-dimensional *unstable* subspace consisting of all $v \in T_x M$ such that

$$\|d\varphi_x^n v\| \leq C \lambda^n \|v\|, \quad (7.3)$$

for all $n \in \mathbb{Z}^-$, and,

- E_x^s is the $(m - 1)$ -dimensional *stable* subspace consisting of all $v \in T_x M$ such that

$$\|d\varphi_x^n v\| \leq C \lambda^n \|v\| \quad (7.4)$$

for all $n \in \mathbb{Z}^+$.

Throughout, we use the shorthand $d\varphi_x$ to write the differential of φ at x , which is a linear map from $T_x M$, the tangent space at x , to $T_{\varphi x} M$, the tangent space at φx . Using the standard basis of \mathbb{R}^m , $d\varphi_x$ can be represented as an $m \times m$ matrix.

Notation 1. For convenience, we write f_x to denote a scalar, vector or tensor field f evaluated at the point $x \in M$.

7.2.5 The SRB measure

An ergodic, φ_s -invariant, physical measure, also known as an SRB measure [281], is guaranteed to exist in our setting of uniformly hyperbolic systems. Apart from φ_s -invariance, we use the ergodicity of the map φ_s with respect to μ_s , which implies that there is no subset of the attractor, other than itself, that is invariant under the dynamics and has full μ_s measure. We also employ the physicality of the SRB measure. This means that ergodic averages starting from Lebesgue-a.e. initial condition chosen in an open set containing the attractor converge to expected values with respect to μ_s .

We also exploit exponential decay of correlations with respect to the SRB measure enjoyed by observables in uniformly hyperbolic systems [74][175][280]. This means that, for two Hölder continuous observables J and f , there is some $c > 0$ and $\delta \in (0, 1)$ so that $|\langle J \circ \varphi^n f \rangle - \langle J \rangle \langle f \rangle| \leq c \delta^n$, for all $n \in \mathbb{Z}^+$. In the S3 algorithm, we often deal with Hölder continuous functions that have zero expectation, in which case, we use $|\langle J \circ \varphi^n f \rangle| \leq c \delta^n$. As a result of exponential decorrelation, Hölder observables also satisfy the central limit theorem (CLT), and the law of the iterated logarithm, which implies that for almost every $x \in M$, the error in the N -time ergodic average of a Hölder observable J declines as $\mathcal{O}(\sqrt{\log \log N}/\sqrt{N})$: $|(1/N) \sum_{n=0}^{N-1} J_{\varphi^n x} - \langle J \rangle| \leq c\sqrt{\log \log N}/\sqrt{N}$, for large N and some $c > 0$. The iterated logarithmic factor, $\sqrt{\log \log N}$, does not significantly alter the convergence rate of ergodic averages from -0.5 , like a Monte Carlo estimation of $\langle J \rangle$. Thus, in this chapter, we write the error convergence of ergodic averages (of Hölder observables) as $\mathcal{O}(1/\sqrt{N})$, where the iterated logarithmic factor is implicit.

Further, we use the fact that the SRB measure, although typically singular with respect to Lebesgue measure on \mathbb{R}^m , has absolutely continuous conditional measures on the unstable manifold. We shall next elaborate on this property as used in the derivation of the S3 algorithm (section 7.3.1).

7.2.6 Parameterization of unstable manifolds

The unstable subspace is tangent to the local unstable manifolds. Given an $\epsilon > 0$, a local unstable manifold at an $x \in \Lambda$, $U_{x,\epsilon}$, contains points whose backward orbits lie ϵ -close to the backward orbit of x . That is,

$$U_{x,\epsilon} = \left\{ x' \in M : \|x_{-n} - x'_{-n}\| \leq \epsilon, \forall n \in \mathbb{Z}^+; \lim_{n \rightarrow \infty} \|x_{-n} - x'_{-n}\| = 0 \right\}.$$

Since E_x^u are one-dimensional subspaces, the local unstable manifolds are also one-dimensional. According to the stable-unstable manifold theorem (see e.g. Theorem 6.2.8 of [154], [89]), the local unstable manifolds are embedded images of Euclidean spaces of the same dimension, which in this case are real lines.

In this chapter, we work with a particular C^1 parameterization of local unstable manifolds. Let Ξ be a measurable partition of Λ , and let Ξ_x denote the element of the partition containing x . At each x , we can choose an ϵ depending on x so that Ξ_x contains a local unstable manifold at every x . We choose a parameterization $\Phi^x : [-\epsilon_x, \epsilon_x] \rightarrow \Xi_x$ that satisfies the following properties:

1. $\Phi^x(0) = x$
2. $d_\xi \Phi^x(\xi) = q_{x'}$, where $\Phi^x(\xi) = x'$, for all x' in the image of Φ^x .

Here $q_{x'} \in T_{x'}M$ is the unit vector in the one-dimensional tangent subspace $E_{x'}^u$. From 2., it follows that $\|d_\xi \Phi^x(\Phi^{x^{-1}}(x'))\| = 1$, for all x' in the image of Φ^x . Hence we refer to the pointwise coordinate maps Φ^x as the unit speed parameterization of local unstable manifolds.

7.2.7 Iterative differentiation on the unstable manifold

The orbits of φ^{-1} starting in Ξ_x generate corresponding orbits on the real line by this parameterization. More concretely, we define the dynamics on the real line, through the map $(\tilde{\varphi}^x)^{-1} := \left(\Phi^{(\varphi^{-1}x)} \right)^{-1} \circ \varphi^{-1} \circ \Phi^x$. We frequently use the following relationship that can be derived using the chain rule, where d_ξ denotes differentiation with respect

to the coordinate ξ :

$$d_\xi (\tilde{\varphi}^x)^{-1} (\xi) = \frac{1}{\alpha_{x'}}, \quad (7.5)$$

Here $\Phi^x(\xi) = x'$, and α is the scalar field that represents the local expansion factor.

Definition 2. *We define a scalar field $\alpha : M \rightarrow \mathbb{R}^+$ to capture the local expansion of unstable tangent vectors. At each $x \in M$,*

$$\alpha_x = \|d\varphi_{\varphi^{-1}x} q_{\varphi^{-1}x}\|. \quad (7.6)$$

Notice that the derivative of $(\tilde{\varphi}^x)^{-1}$ with respect to the unstable coordinate ξ does not depend on the base point of the coordinate system, x . This fact is crucial to the derivation of the S3 algorithm, where we prominently differentiate scalar and vector fields “on unstable manifolds” of a reference orbit. We now describe what this differentiation means in this chapter, and a formula useful for performing this differentiation recursively.

Let $x_0 \in M$ be a fixed μ -typical point whose forward orbit serves as the reference orbit for our computation. Let $f : M \rightarrow \mathbb{R}$ be a scalar field, and h be the scalar field that represents the derivative of f on the unstable manifold. We define this function, using our coordinate systems centered on $\{x_n\}$ as follows, $h_{x'_n} := d_\xi(f \circ \Phi^{x_n})((\Phi^{x_n})^{-1}(x'_n))$, for some $x'_n \in \Xi_{x_n}$.

Now suppose we wish to compute h recursively along an orbit; the situation where such a computation arises in the S3 algorithm is discussed in section 7.5. To compute the values h_{x_n} using $\{h_{x_m}\}_{m < n}$, we use the following iteration that in turn uses Eq. 7.5,

$$\begin{aligned} d_\xi(f \circ \varphi \circ \Phi^x)(0) &= d_\xi(f \circ \Phi^{\varphi x} \circ \tilde{\varphi}^{\varphi x})(0) \\ &= h_{\varphi x} \alpha_{\varphi x}. \end{aligned} \quad (7.7)$$

We sometimes use the notation $\partial_\xi f$ to denote h , which we simply refer to as the unstable derivative of f .

7.2.8 Conditional density of the SRB measure on unstable manifolds

An important property of the SRB measure, which guarantees its physicality, is its absolute continuity along unstable manifolds [281]. To define this property, we must introduce conditional densities of the SRB measure on the unstable manifold, which we denote ρ . We may define the function ρ using our parameterization $\{\Phi^x\}$ and disintegration to yield a family of functions $\{\rho^x\}$. It is worth noting that changes to the parameterization only affect the normalization constant of the densities ρ^x . By the disintegration of SRB measure [169][76] on the measurable partition Ξ , there exist i) a family of conditional measures denoted μ^x , at μ -a.e. $x \in M$, that are supported on Ξ_x , and ii) a quotient measure $\hat{\mu}$ on M/Ξ , such that, for all Borel subsets $A \subseteq M$,

$$\mu(A) = \int_{M/\Xi} \mu^x(A \cap \Xi_x) d\hat{\mu}(x). \quad (7.8)$$

An aside on the notation used henceforth: a phase point $x \in M$ that appears in the superscript indicates the point at which our pointwise coordinate system is centered; this is distinct from its appearance on a subscript, which always means evaluation of a scalar/vector field at x . Let ν^x be the normalized pushforward of the Lebesgue measure (uniform probability distribution) on $[0, 1]$, by Φ^x . The absolute continuity property of the SRB measure on the unstable manifold means that the conditional measure μ^x is absolutely continuous with respect to ν^x . Thus, a scalar function ρ^x can be defined as the *probability density* of μ^x . In particular, the unnormalized density ρ_u^x of μ^x has been derived by Pesin [218][148] to be,

$$\rho_u^x(x') := \prod_{k=0}^{\infty} \frac{\alpha_{\varphi^{-k}x}}{\alpha_{\varphi^{-k}x'}}, \quad (7.9)$$

where the local expansion factor α is as defined in Eq. 7.6.

Definition 3. *The probability density of conditional SRB measures on the unstable manifold is defined as $\rho^x(x') = \rho_u^x(x')/\bar{\rho}^x$, for $x' \in \Xi_x$, where $\bar{\rho}^x := \int_{\Xi_x} \rho_u^x(x') d\nu^x(x')$, is the normalization constant.*

Using the above definition of the density, the disintegration of the SRB measure on the unstable manifold (Eq. 7.8) gives, for any smooth observable f ,

$$\langle f \rangle = \int_{M/\Xi} \int_{\Xi_x} f(x') \rho(x') d\nu^x(x') d\hat{\mu}(x). \quad (7.10)$$

7.2.9 Linear response of uniformly hyperbolic systems

Ruelle [232][234] showed that linear response holds and is given by the following formula in uniformly hyperbolic attractors,

$$\langle J, \partial_s \mu_s \rangle = \sum_{k=0}^{\infty} \langle d(J \circ \varphi^k) \cdot \chi, \mu \rangle, \quad (7.11)$$

where, as usual, we have dropped the subscript on φ and μ to indicate their respective values at $s = s_0$. According to Ruelle's formula (7.11), linear response or the parametric derivative of statistics, is a series summation of ensemble averages. The sensitivity of the function $J \circ \varphi^k$ to the parameter perturbation χ is given by $d(J \circ \varphi^k) \cdot \chi$, and the k th term in Ruelle's formula is an ensemble average of this instantaneous sensitivity. With probability 1, the instantaneous sensitivity, $d(J \circ \varphi^k) \cdot \chi$, grows in norm exponentially with k since χ almost surely has a non-zero component in E^u . But, due to cancellations over phase space, the ensemble average is bounded at all k . Further, the convergence of Ruelle's series implies that the ensemble average of instantaneous sensitivities decreases asymptotically with time and converges to zero.

Numerically, the direct evaluation of Ruelle's formula (7.11) involves approximating each ensemble average (each term in the series) as a sample average of instantaneous sensitivities. These instantaneous sensitivities are, in turn, computed by using the conventional tangent equation (section 7.2.2) or when the dimension of the parameter space is large, using the adjoint equation [211][96][69]. One may also use automatic differentiation in the forward and reverse mode to approximate the tangent and adjoint solutions respectively. However, the resulting sensitivity from any linear perturbation method increases exponentially with k at the rate of the largest Lyapunov exponent, λ_1 ; hence considering the sample average of the sensitivities as

a random variable, its variance increases as $\mathcal{O}(e^{2\lambda_1 k})$. Thus, the number of samples needed to approximate the integral accurately rapidly increases with k . On the other hand, thresholding the series computation (Eq. 7.11) at a small value of k to reduce the variance introduces a bias. This bias-variance trade-off has been analyzed in previous works [96][69], and the direct evaluation of Ruelle's formula rendered infeasible in practical systems.

7.3 Main results: S3 decomposition and computation of Ruelle's response formula

The two key ideas that enable an efficient, rigorous computation of Ruelle's formula are i) a particular decomposition of the perturbation field χ that splits Ruelle's formula into *stable* and *unstable* contributions, and ii) the development of iterative, ergodic-averaging evaluations of these two components. The first main result concerns the specific decomposition of a given parameter perturbation that we show enables an efficient computation.

Theorem 2. *A differentiable vector field χ has a sequence of decompositions $\chi = a^n q + (\chi - a^n q)$, where q is the unit vector field tangent to the one-dimensional unstable manifold, such that*

1. *the sequence of vector fields, $\{v^n\}_{n \in \mathbb{Z}^+}$ that satisfies*

$$v_x^{n+1} := d\varphi_{\varphi^{-1}x} v_{\varphi^{-1}x}^n + (\chi_x - a_x^{n+1} q_x), \quad x \in M, n \in \mathbb{Z}^+, \quad (7.12)$$

with v^0 being any bounded vector field that is differentiable on the unstable manifold, converges uniformly; the sequence of scalar fields $\{a^n\}$ is chosen to orthogonalize v^n to the unstable manifold: $v_x^n \cdot q_x = 0$, at all $x \in M$ and for all $n \in \mathbb{Z}^+$;

2. *the sequence $\{a^n\}$ is differentiable on the unstable manifold and uniformly converging; and,*

3. the sequence of unstable derivatives of $\{a^n\}$, namely, $\{b^n := \partial_\xi a^n\}$ converges uniformly.

Let the limits of the sequences we introduced in Theorem 2 be denoted as follows: $a := \lim_{n \rightarrow \infty} a^n$, $v := \lim_{n \rightarrow \infty} v^n$ and $b := \lim_{n \rightarrow \infty} b^n$. From Theorem 2-2. and 2-3., the scalar field b is the unstable derivative of a , i.e., $b = \partial_\xi a$. We use the sequence $\{v^n\}$ to split Ruelle's formula into two computable infinite series. Fixing some finite K that controls the accuracy of the implementation, the S3 decomposition can be construed as a sequence of decompositions of χ into $\chi = (a^{K-k} q) + (\chi - a^{K-k} q)$, $0 \leq k < K$, wherein the former component is tangent to the unstable manifold. This gives the following decomposition, which converges to Ruelle's formula, as $K \rightarrow \infty$,

$$\langle J, \partial_s \mu_s \rangle = \lim_{K \rightarrow \infty} \sum_{k=0}^{K-1} \langle d(J \circ \varphi^k) \cdot (\chi - a^{K-k} q), \mu \rangle + \lim_{K \rightarrow \infty} \sum_{k=0}^{K-1} \langle d(J \circ \varphi^k) \cdot a^{K-k} q, \mu \rangle. \quad (7.13)$$

We refer to the first term on the right hand side as the *stable contribution*, and the second as the *unstable contribution* to the overall sensitivity. In section 7.4, we use Theorem 2 to show that the stable and unstable contributions can be alternatively expressed in the following forms that are amenable to their computation.

$$\langle J, \partial_s \mu_s \rangle = \langle dJ \cdot v, \mu \rangle - \lim_{K \rightarrow \infty} \sum_{k=0}^{K-1} \langle J \circ \varphi^k (a^{K-k} g + b^{K-k}), \mu \rangle, \quad (7.14)$$

where we have introduced the *logarithmic density gradient* g , which is defined, for $x \in \Xi_{x'}$, as

$$g_x := \frac{1}{\rho^{x'}(x)} \frac{d(\rho^{x'} \circ \Phi^{x'})}{d\xi}((\Phi^{x'})^{-1}(x)). \quad (7.15)$$

The intuition for the S3 decomposition is developed in section 7.4, and the proof of Theorem 2 is presented in section 7.7.

7.3.1 The S3 algorithm

The S3 algorithm is an efficient, rigorous computation of the above split Ruelle's formula (Eq. 7.14). The computation of the different ingredients of this new formula as an efficient ergodic average is the other major contribution of this chapter. In particular, we develop iterative algorithms to compute v , g and b along trajectories. Before we detail the S3 algorithm, we recall an important notation used throughout this chapter.

Notation 2. *When a μ -typical phase point $x_0 \in M$ is fixed, the subscript n applied to a scalar function or a vector field, denoted f , refers to its corresponding value at x_n . That is, when $x_0 \in M$ is fixed, $f_n := f_{\varphi^n x_0} = f_{x_n}$.*

The S3 algorithm is as follows:

1. Obtain a long primal trajectory $x_{-K'}, \dots, x_{N-1}$, where $x_{n+1} = \varphi x_n$, $-K' \leq n \leq (N-1)$, with $x_{-K'}$ chosen μ -a.e.
2. Obtain, at each point x_n , the unit tangent vector to the unstable manifold, q_n . The following procedure converges exponentially in n to the true value of q_n . Solve the homogeneous tangent equation with repeated normalization. That is, solve

$$\alpha_{n+1} q_{n+1} = (d\varphi)_n q_n, \quad n = -K', \dots, 0, 1, \dots, \quad (7.16)$$

with q_0 being set to a random vector in \mathbb{R}^m , and $\alpha_{n+1} = \|(d\varphi)_n q_n\|$.

3. Solve for v_n the following inhomogeneous tangent equation, which repeatedly projects v_n out of the unstable subspace,

$$v_{n+1} = (d\varphi)_n v_n + \chi_{n+1} - a_{n+1} q_{n+1}, \quad n = -K', \dots, 0, 1, \dots, \quad (7.17)$$

where a_{n+1} is such that $v_{n+1} \cdot q_{n+1} = 0$. The initial condition $v_{-K'}$ is set to $0 \in \mathbb{R}^m$. This equation is henceforth called the *regularized* tangent equation.

The error in the solutions v_n when compared to the true value of the vector field v (Theorem 2) decreases exponentially with n (as shown in Lemma 3).

4. Solve the following second-order tangent equation for p_n , starting with $p_{-K'} = 0 \in \mathbb{R}^m$,

$$p_{n+1} = \frac{(d^2\varphi)_n(q_n, q_n) + (d\varphi)_n p_n}{\alpha_{n+1}^2}, \quad n = -K', \dots, 0, 1, \dots, \quad (7.18)$$

The solutions p_n converge exponentially with n to the true values of the vector field p along the orbit (Lemma 4).

5. Solve the following recursive second-order tangent equation for y_n , for each $n \in \mathbb{Z}^+$,

$$y_{n+1} = \frac{(d^2\varphi)_n(q_n, v_n) + (d\varphi)_n y_n}{\alpha_{n+1}} + (d\chi)_{n+1} q_{n+1} - c_{n+1} q_{n+1} - a_{n+1} p_{n+1}, \quad n = -K', \dots, 0, \dots, \quad (7.19)$$

where the scalar c_{n+1} is found using the relationship $y_{n+1} \cdot q_{n+1} = -v_{n+1} \cdot p_{n+1}$. This relationship follows from taking the unstable derivative of $v_{n+1} \cdot q_{n+1} = 0$. The y_n and c_n computed using this procedure also converge exponentially with n to their true values, as shown in Lemma 6 and 7 respectively.

6. Compute the unstable contribution as the following ergodic average:

$$\langle J, \partial_s \mu_s \rangle^u \approx -\frac{1}{N} \sum_{k=0}^{K-1} \sum_{n=0}^{N-1} J_{n+k} c_n. \quad (7.20)$$

7. Compute the stable contribution as the following ergodic average:

$$\langle J, \partial_s \mu_s \rangle^s \approx \frac{1}{N} \sum_{n=0}^{N-1} (dJ)_n \cdot v_n. \quad (7.21)$$

8. The output of the S3 algorithm is the sum of the stable and unstable contributions, $\langle J, \partial_s \mu_s \rangle^u + \langle J, \partial_s \mu_s \rangle^s$.

The following theorem establishes the convergence of the above S3 algorithm.

Theorem 3. *The ergodic averages in Eq. 7.20 and Eq. 7.21 converge to the unstable and stable contributions respectively. In particular,*

1. *for μ -a.e. x_0 and for almost every q_0 ,*

$$\langle J, \partial_s \mu_s \rangle^u = - \lim_{K \rightarrow \infty} \lim_{N \rightarrow \infty} \sum_{k=0}^{K-1} \frac{1}{N} \sum_{n=0}^{N-1} J_{n+k} c_n; \quad (7.22)$$

2. *And, for μ -a.e. x_0 ,*

$$\langle J, \partial_s \mu_s \rangle^s = \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{n=0}^{N-1} (dJ)_n \cdot v_n. \quad (7.23)$$

We prove Theorem 3 in section 7.8. The efficiency of the S3 algorithm in comparison to a direct evaluation of Ruelle's formula (Eq. 7.11) stems from the following. The integrand in Ruelle's original series increases in norm exponentially with k . This makes the ergodic averaging approximation of the integral computationally inefficient, as we noted in section 7.2.9. By contrast, the norm of the integrand in the S3 modification of the formula (Eq. 7.14) is uniformly bounded in both the stable and unstable contributions. Thus, efficient ergodic-averaging approximations are possible for the S3 formula (Eq. 7.14).

7.4 Derivation and computation of the S3 formula

In this section and the next, we describe the S3 decomposition of Ruelle's formula (Eq. 7.13), the derivation of its regularized form (Eq. 7.14), and the computation of the latter. In order to derive an efficient computation of Ruelle's formula, we split Ruelle's formula into two parts, using a particular decomposition of the vector field. The contribution to the overall sensitivity made by the unstable component, which is aligned with q , is called the unstable contribution; the remaining term is the stable contribution (Eq. 7.13).

However, we do not decompose χ into its stable and unstable components. That is, although $aq \in E^u$ clearly, a is not chosen such that $\chi - aq$ belongs to E^s . Instead, the scalar field a is chosen so as to be differentiable on the unstable manifold and such that both the stable and unstable contributions lead to well-conditioned computations. The significance of the differentiability of a on the unstable manifold will be clear in the derivation below.

7.4.1 Regularizing tangent equation solutions

The S3 decomposition can be motivated as a means of regularizing the solutions of a conventional tangent equation. Fixing a reference orbit $\{x_n\}_{n \in \mathbb{Z}^+}$, consider the conventional tangent equation (Eq. 7.2), which is a recursive equation for $u_n := (\partial_s \varphi_s^n)(x_0, s)$, starting from $u_0 = 0 \in \mathbb{R}^m$. As noted in section 7.2.2, this equation gives the evolution of a tangent vector corresponding to the parameter perturbation along a fixed trajectory, $\{x_n\}_{n \in \mathbb{Z}^+}$. By construction (Eq. 7.2), we can see that $u_n = \sum_{k=0}^{n-1} (d\varphi^k)_{n-k} \chi_{n-k}$. By definition of chaos (section 7.2.3), $\|d\varphi_x^k \chi_x\| \sim \mathcal{O}(e^{\lambda_1 k})$, for almost every $x \in M$, and almost every perturbation $\chi_x \in \mathbb{R}^m$. Hence, for large n , $\|u_n\| \sim \mathcal{O}(e^{\lambda_1 n})$ for almost every perturbation χ . On the other hand, if we projected out the unstable component of the tangent solution at each timestep, the solution does not exhibit exponential growth. That is, the following iteration is stable

$$v_{n+1} = (d\varphi)_n v_n + (\chi_{n+1} - a_{n+1} q_{n+1}), \quad (7.24)$$

where a_{n+1} is such that $v_{n+1} \cdot q_{n+1} = 0$, for all $n \in \mathbb{Z}^+$. That is, the solution v_n of the above tangent equation along with the repeated projections out of the unstable subspace, is in $E_{x_n}^{u\perp}$ at each n . We refer to the solutions $\{v_n\}$ as *regularized tangent solutions*. With this stable iterative procedure as the motivation, we derive a splitting of Ruelle's formula. One part of the split formula – the stable contribution – will be computed using the regularized tangent solution, Eq. 7.24.

7.4.2 Alternative expression of the stable contribution

In the regularized tangent equation (Eq. 7.24), a tangent vector is projected out of the unstable subspace at every iteration. At every $x \in M$, we introduce the orthogonal projection operator, \mathcal{P}_x , to refer to this operation.

Notation 3. Let $\mathcal{P}_x : T_x M \rightarrow T_x M$ denote the linear operator that projects a tangent vector onto the hyperplane orthogonal to E_x^u . If q_x is a unit vector in E_x^u , then for any $v_x \in T_x M$, $\mathcal{P}_x v_x := (I - q_x q_x^T) v_x$. Applying \mathcal{P}_x to every point on M would result in a linear operator, hereafter denoted \mathcal{P} , on vector fields of M .

Using this notation, we can reproduce the operations performed by solving Eq. 7.24, by considering a sequence of vector fields $\{v^n\}$ that satisfies the following recurrence relation,

$$v^{n+1} = \mathcal{P}(d\varphi v^n + \chi), \quad n \in \mathbb{Z}^+. \quad (7.25)$$

We also introduce a sequence of scalar fields which denote the components on the unstable subspace that are projected out every iteration. That is, let $\{a^n\}$ be a sequence of scalar fields given by,

$$a^{n+1} := q^T(d\varphi v^n + \chi). \quad (7.26)$$

In Lemma 3, we show that the sequence $\{v^n\}$ converges uniformly, starting from any bounded vector field v^0 ; the limit of this sequence is denoted v . The uniform convergence of $\{v^n\}$ implies the uniform convergence of the scalar field $\{a^n\}$, as we show in Lemma 5; the limit of this sequence is denoted a . Using these results and uniform hyperbolicity, we prove the following alternative formula for the stable contribution in Proposition 3:

$$\begin{aligned} \langle J, \partial_s \mu_s \rangle^s &:= \lim_{K \rightarrow \infty} \sum_{k=0}^{K-1} \langle d(J \circ \varphi^k) \cdot (\chi - a^{K-k} q), \mu \rangle \\ &= \langle dJ \cdot v, \mu_s \rangle. \end{aligned} \quad (7.27)$$

7.4.3 Computation of the stable contribution

That is, the series summation representing the stable contribution is simply an ensemble average of $dJ \cdot v$, an inner product of two bounded vector fields. That is, for x_0 chosen μ -a.e.,

$$\langle J, \partial_s \mu_s \rangle = \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{n=0}^{N-1} (dJ)_n \cdot v_n. \quad (7.28)$$

The values of the vector field v along the reference orbit are approached by the regularized tangent solution (Eq. 7.24). Without loss of generality, we effect starting the recursive Eq. 7.25 from a zero vector field, by choosing $v_{-K'} = 0 \in \mathbb{R}^m$, at the point $x_{-K'}$. For a large run-up time K' , the solution at 0, v_0 and a_0 are already close to the true values of v and a at x_0 . The solutions thus produced by long-time evolution of the regularized tangent equation (Eq. 7.24) become exponentially more accurate, and are used to evaluate the stable contribution as per Eq. 7.28. In practice, we solve the regularized tangent equation (Eq. 7.24) over an orbit of a long but finite length N . This computes an approximation of the limit on the right hand side up to a finite N . Proposition 3 shows that such a computation of the stable contribution converges to its true value (Eq. 7.27) as $N \rightarrow \infty$; the asymptotic error convergence is as $\mathcal{O}(1/\sqrt{N})$, ignoring log factors.

7.4.4 Alternative expression of the unstable contribution

Recall that the unstable contribution is the sensitivity due to the perturbation along the unstable subspace,

$$\langle J, \partial_s \mu_s \rangle^u = \lim_{K \rightarrow \infty} \sum_{k=0}^{K-1} \langle d(J \circ \varphi^k) \cdot a^{K-k} q, \mu \rangle. \quad (7.29)$$

As noted in section 7.2.9, the integrand in the k th term of the above series, $d(J \circ \varphi^k) \cdot a^{K-k} q$, increases exponentially in norm with k . Thus, rather than a direct evaluation, we apply integration by parts on the unstable manifold, which moves the derivative

away from the time-dependent term $J \circ \varphi^k$.

Before we integrate by parts, we apply disintegration (Eq. 7.10) of the SRB measure to rewrite the unstable contribution,

$$\begin{aligned} \langle J, \partial_s \mu_s \rangle^u &:= \lim_{K \rightarrow \infty} \sum_{k=0}^{K-1} \langle a^{K-k} d(J \circ \varphi^k) \cdot q, \mu \rangle \\ &= \lim_{K \rightarrow \infty} \sum_{k=0}^{K-1} \int_{M/\Xi} \int_{B_x} a^{K-k} \circ \Phi^x(\xi) \frac{d(J \circ \varphi^k \circ \Phi^x)}{d\xi}(\xi) \rho^x \circ \Phi^x(\xi) d\xi d\hat{\mu}(x), \end{aligned} \quad (7.30)$$

where $B_x := \Phi^{x-1}(\Xi_x)$. When compared to Eq. 7.10, we have additionally used a change of variables, $x' \rightarrow \xi$, in the inner integral. Since $\|\Phi^{x'}(\xi)\| = 1$, by construction, this change of variables does not introduce a multiplicative factor in the integrand. Now applying integration by parts on the inner integral, the k th term on the right hand side of Eq. 7.30 becomes,

$$\begin{aligned} &\int_{M/\Xi} \int_{\partial B_x} \frac{d((a^{K-k} J \circ \varphi^k \rho^x) \circ \Phi^x)}{d\xi} d\xi d\hat{\mu}(x) \\ &- \int_{M/\Xi} \int_{B_x} J \circ \varphi^k \circ \Phi^x \left(\frac{d(a^{K-k} \circ \Phi^x)}{d\xi} + \frac{a^{K-k} \circ \Phi^x d(\rho^x \circ \Phi^x)}{\rho^x \circ \Phi^x d\xi} \right) (\rho^x \circ \Phi^x) d\xi d\hat{\mu}(x). \end{aligned} \quad (7.31)$$

The first term in the above equation vanishes due to cancellations on the boundaries of B_x for different x (Theorem 3.1(b) in [232][234]; see also [148]). Changing variables back to x' , we obtain the following regularized expression for the unstable contribution,

$$\begin{aligned} \langle J, \partial_s \mu_s \rangle^u &= - \lim_{K \rightarrow \infty} \sum_{k < K} \int_{M/\Xi} \int_{\Xi_x} J \circ \varphi^k(x') \left(\frac{d(a^{K-k} \circ \Phi^x)}{d\xi} ((\Phi^x)^{-1}(x')) \right. \\ &\quad \left. + \frac{a^{K-k}(x') d(\rho^x \circ \Phi^x)}{\rho^x(x') d\xi} ((\Phi^x)^{-1}(x')) \right) \rho^x(x') d\nu_x(x') d\hat{\mu}(x), \end{aligned} \quad (7.32)$$

We introduce the logarithmic density gradient function [247][210],

$$g^x(x') := \frac{1}{\rho^x(x')} \frac{d(\rho^x \circ \Phi^x)}{d\xi} ((\Phi^x)^{-1}(x')). \quad (7.33)$$

As we show in section 7.5.2, g^x does not depend on x , and hence we denote the density gradient simply as g . We also introduce the scalar field sequence b^k defined as $b_x^k := d_\xi(a^k \circ \Phi^x)(0)$ to denote the unstable derivative of a^k , deferring the (constructive) proof of its uniform convergence until Lemma 7. This leads us to the following expression for the unstable contribution,

$$\langle J, \partial_s \mu_s \rangle^u = - \lim_{K \rightarrow \infty} \sum_{k=0}^{K-1} \langle J \circ \varphi^k (g a^{K-k} + b^{K-k}), \mu \rangle. \quad (7.34)$$

Recall that we compute a^k iteratively along a typical trajectory as part of the stable contribution computation (section 7.4.2). We are now left with the computation of g and b^k along the trajectory. These are tackled in section 7.5.2 and 7.5.3 respectively. In a recent work [210], Ni addresses the computation of the divergence on the unstable manifold (analogous to the function g described above), where the differentiation is performed on a set of *shadowing* coordinates that are tied to the parameter perturbation. This is part of an algorithm to compute linear response that uses a different decomposition from the present chapter, into shadowing and unstable directions.

7.5 Computing derivatives along unstable manifolds

In the previous section, we obtained a regularized expression for the unstable contribution. That is, the integrand in the regularized expression, Eq. 7.34, is uniformly bounded – the uniform boundedness of the scalar field sequences $\{a^k\}$ and $\{b^k\}$ and the boundedness of g are proved in Lemma 5, 7 and 8 respectively. The question still remains how we can compute the unknown scalar fields g and b . For both these computations, we require second-order unstable derivatives, which we discuss next.

7.5.1 Iteratively computing the curvature of the unstable manifold

The function $\Phi^{x'}(\xi)$ is an arclength travelled by a particle on a local unstable manifold of x' , starting from x' . The particle travels with unit speed, and its instantaneous velocity at time ξ is $q_{\Phi^{x'}(\xi)}$. Its acceleration field is given by $w^{x'} := d_\xi^2 \Phi^{x'}(\xi)$. We derive a recursive equation satisfied by this family, by starting with a differentiation with respect to ξ of the definition of α (Eq. 7.6),

$$\alpha_{\varphi x}^2 \gamma_{\varphi x}^{\varphi x'} = (d\varphi_x q_x)^T (d\varphi_x w_x^{x'} + d^2\varphi_x(q_x, q_x)). \quad (7.35)$$

Here $\gamma_x^{x'} := d_\xi(\alpha \circ \Phi^{x'})((\Phi^{x'})^{-1}x)$ is the unstable derivative of α ; $d^2\varphi_x$ is the bilinear form that returns a tangent vector $\in T_{\varphi x}M$; it can be written as an $m \times m \times m$ tensor consisting of component-wise partial derivatives of the Jacobian $d\varphi_x$. In deriving Eq. 7.35, we have used the chain rule in Eq. 7.7. Now differentiating the statement of covariance of the unstable subspace (section 7.2.4), $\alpha_{\varphi x} q_{\varphi x} = d\varphi_x q_x$,

$$\alpha_{\varphi x} \gamma_{\varphi x}^{x'} q_{\varphi x} + \alpha_{\varphi x}^2 w_{\varphi x}^{\varphi x'} = d\varphi_x w_x^{x'} + d^2\varphi_x(q_x, q_x). \quad (7.36)$$

Substituting Eq. 7.35 into Eq. 7.36, and using the definition of α ,

$$w_{\varphi x}^{\varphi x'} = \frac{1}{\alpha_{\varphi x}^2} (I - q_{\varphi x} q_{\varphi x}^T) \left(d\varphi_x w_x^{x'} + d^2\varphi_x(q_x, q_x) \right), \quad (7.37)$$

where I is the $m \times m$ identity matrix. The above equation represents the following relationship between elements of the family of vector fields, $\{w^x\}$,

$$w^{\varphi x} = \mathcal{P} \frac{d\varphi w^x + d^2\varphi(q, q)}{\alpha^2}. \quad (7.38)$$

In Lemma 4, we show that any bounded sequence $\{w^n\}$ of vector fields that satisfies,

$$w^{n+1} = \mathcal{P} \frac{d\varphi w^n + d^2\varphi(q, q)}{\alpha^2} \quad (7.39)$$

converges uniformly to a unique vector field w . Hence, the family $\{w^x\}$ does not depend on the parameterization centers, and simply denotes a single vector field, which we call w . Reconsidering Eq. 7.35,

$$\gamma_{\varphi x}^{\varphi x'} = \frac{q_{\varphi x}^T}{\alpha_{\varphi x}} (d\varphi_x w_x + d^2\varphi_x(q_x, q_x)), \quad (7.40)$$

it is clear that the family of scalar fields $\{\gamma^x\}$ is also independent of the parameterization. We henceforth write $\gamma = \partial_\xi \alpha$, to denote the unstable derivative of the scalar field α .

Similarly, the scalar fields $d_\xi(a \circ \Phi^x)$ and $d_\xi(\log \rho^x \circ \Phi^x)$ can be attained as limits of (exponentially) uniformly converging sequences of scalar fields, as proved in Lemma 7 and Lemma 8 respectively. Thus, by the same argument as we used for $\{w^x\}$, these limits are independent of the parameterization centers and are denoted simply $b := \partial_\xi a$ and $g := \partial_\xi \log \rho^x$ respectively.

The curvature of a local unstable manifold is the magnitude of the acceleration experienced by a particle travelling at unit speed. That is, the curvature of the unstable manifold at x is $\|w_x\|$. From Eq. 7.35, it is clear that the computation of the curvature and that of γ go hand in hand. To compute both w and then using Eq. 7.35, γ , along a typical trajectory, we solve the second-order tangent equation derived above (Eq. 7.39) iteratively along the trajectory. In practice, we assume $w_0^0 := w_{x_0}^0 = 0$ at some μ -typical x_0 and iterate Eq. 7.39. Such a computation converges exponentially with n to the true value of w_n due to Lemma 4. At each step of the recursion, the value of γ_n , which is computed through Eq. 7.35 using the computed values of w_n , also becomes exponentially more accurate. These values of γ_n along the orbit $\{x_n\}$ are used to obtain g along the orbit, as we shall discuss in the next subsection.

From Eq. 7.39, we can also infer that the vector field w is orthogonal to q (i.e., the acceleration of a particle on a local unstable manifold is perpendicular to its velocity). Visualizations of the unstable manifold curvatures obtained from the norms of w_n computed as above are shown for classical hyperbolic attractors in [63].

7.5.2 Iterative formula for unstable derivatives of SRB density

In a previous work [247], we provide an intuitive explanation of the density gradient g as well as its computation in one-dimensional expanding maps of the interval, where the global unstable manifold is the background manifold M . In the present setting of a one-dimensional unstable manifold, we see that our derivation leads to a similar computation of g as in 1D expanding maps. From our expression for the SRB density (Eq. 3), we find that for $x \in \Xi_{x'}$,

$$\rho_x^{x'} = \rho_{\varphi^{-1}x}^{\varphi^{-1}x'} \frac{\alpha_{x'}}{\alpha_x} \frac{\bar{\rho}^{\varphi^{-1}x'}}{\bar{\rho}^{x'}}. \quad (7.41)$$

We recall that phase points appearing on superscripts indicate the centers of our coordinate system Φ^x (section 7.2.6), while on subscripts, they indicate evaluations of the function (e.g., in Eq. 7.41, the scalar functions $\rho^{x'}$ and $\rho^{\varphi^{-1}x'}$ are evaluated at x and $\varphi^{-1}x$ respectively). Taking logarithm and differentiating Eq. 7.41 with respect to ξ on the unstable manifold at x , using the definition (Eq. 7.33) of the density gradient g^x ,

$$\begin{aligned} g_x^{x'} &= \frac{g_{\varphi^{-1}x}^{\varphi^{-1}x'}}{\alpha_x} - \frac{1}{\alpha_x} \frac{d(\alpha \circ \Phi^{x'})}{d\xi} ((\Phi^{x'})^{-1}(x)) \\ &= \frac{g_{\varphi^{-1}x}^{\varphi^{-1}x'}}{\alpha_x} - \frac{\gamma_x}{\alpha_x}. \end{aligned} \quad (7.42)$$

To derive the first term on the right hand side, we have also used the chain rule for our parameterization (Eq. 7.7) and the scalar field γ introduced in the previous subsection. Now, in Lemma 8, we show that starting from any bounded scalar function h^0 , and considering any bounded function r , the following recursion converges uniformly,

$$h_{\varphi x}^{n+1} = \frac{h_x^n}{\alpha_{\varphi x}} + r_{\varphi x}. \quad (7.43)$$

Since the sequence of functions $\{g^{\varphi^n x}\}_{n=-\infty}^0$ in Eq. 7.42 satisfies this same recursion, the family of functions $\{g^x\}$ must indeed be a single function independent of x . Hence,

we omit the superscript x , and simply denote the density gradient g . Thus, Eq. 7.42 can be rewritten as follows, fixing some μ -typical reference orbit $\{x_n\}$:

$$g_{n+1} = \frac{g_n}{\alpha_{n+1}} - \frac{\gamma_{n+1}}{\alpha_{n+1}}. \quad (7.44)$$

This iterative equation exponentially converges (Lemma 8), with n , to the true value of $g_n := g(x_n)$, along almost every orbit. In practice, we start the computation with an arbitrary initialization.

7.5.3 Iterative computation of the scalar field b

Having completely prescribed a computation for g , the only unknown we are left with in computing the unstable contribution (Eq. 7.34) is the scalar field b . We describe a procedure, involving w , and a new second-order tangent equation. This latter second-order tangent solution exponentially approaches the vector field representing the unstable derivative of v . In order to derive this equation, recall that the vector field v is the limit of a sequence $\{v^n\}$ described in Eq. 7.25. The sequence $\{v^n\}$ is differentiable in the unstable direction, if we start with a differentiable vector field v^0 . In Lemma 6, we show that the sequence of these derivatives, denoted $\{y^n\}$, converges uniformly. Since $\{v^n\}$ also converges uniformly, the limit y of the sequence $\{y^n\}$ is the unstable derivative of v .

Thus, it is valid to differentiate the regularized tangent equation (Eq. 7.24) in the unstable direction. Taking this derivative, we obtain that the vector field y satisfies the following equation, along a fixed μ -typical orbit $\{x_n\}_{n \in \mathbb{Z}^+}$

$$\begin{aligned} \alpha_{n+1}y_{n+1} &= (d\varphi)_n y_n + (d^2\varphi)_n(q_n, v_n) \\ &+ \alpha_{n+1}(d\chi)_{n+1} q_{n+1} - \alpha_{n+1}b_{n+1}q_{n+1} - \alpha_{n+1}a_{n+1}w_{n+1}. \end{aligned} \quad (7.45)$$

As usual, we have used the chain rule for differentiating on unstable manifolds (Eq. 7.7) and the differentiability of $d\varphi$ and χ on M . We have suppressed the parameterization centers on the superscripts because the unstable derivatives y and b are

limits of uniformly converging series, and hence we can invoke the same argument as in sections 7.5.1 and 7.5.2 to show their independence from the parameterization. Using Lemma 6, we can argue that the iteration above converges to the true value of the vector y_n as $n \rightarrow \infty$, starting with an arbitrary value for y_0 . In order to compute the above recursion, we must know the value of b along the trajectory. However, it is possible to close this system of equations for y_{n+1} and b_{n+1} by differentiating in the unstable direction, the definition of a_{n+1} .

Recall that the regularized tangent solutions v_n are orthogonal to q_n and a_n are chosen so as to enforce this orthogonality. From Eq. 7.24, we get,

$$a_{n+1} = q_{n+1} \cdot ((d\varphi)_n v_n + \chi_{n+1}). \quad (7.46)$$

Differentiating the above equation in the unstable direction,

$$\begin{aligned} \alpha_{n+1} b_{n+1} &= ((d^2\varphi)_n(v_n, q_n) + (d\varphi)_n y_n + \alpha_{n+1}(d\chi_{n+1} q_{n+1})) \cdot q_{n+1} \\ &+ \alpha_{n+1} w_{n+1} \cdot ((d\varphi)_n v_n + \chi_{n+1}). \end{aligned} \quad (7.47)$$

By comparing Eq. 7.47 with the inner product of Eq. 7.45 with q_{n+1} , we see that

$$y_{n+1} \cdot q_{n+1} = -w_{n+1} \cdot ((d\varphi)_n v_n + \chi_{n+1}). \quad (7.48)$$

By using the orthogonality $w_{n+1} \cdot q_{n+1} = 0$, which follows from Eq. 7.37, on Eq. 7.24, we obtain $w_{n+1} \cdot v_{n+1} = w_{n+1} \cdot ((d\varphi)_n v_n + \chi_{n+1})$. Thus, we can further simplify Eq. 7.48 to obtain the constraint,

$$y_{n+1} \cdot q_{n+1} = -w_{n+1} \cdot v_{n+1}. \quad (7.49)$$

Thus, we have two equations (Eq. 7.45 and Eq. 7.49) from which the two unknowns b_{n+1} and y_{n+1} are solved for. Without loss of generality, we start the iteration in Eq. 7.45 with $y_0 = 0 \in \mathbb{R}^m$. Then, by Lemma 6 and Lemma 7, this iterative procedure yields exponentially accurate values of both y_n and b_n along a trajectory.

7.5.4 Computation of the unstable contribution

We now assemble the components computed in the previous subsections together to form the unstable contribution, which we recall from Eq. 7.34,

$$\langle J, \partial_s \mu_s \rangle^u = - \lim_{K \rightarrow \infty} \sum_{k=0}^{K-1} \langle J \circ \varphi^k (a^{K-k} g + b^{K-k}), \mu \rangle. \quad (7.50)$$

We shall assume a strong form of exponential correlation decay (section 7.2.5) by which, for a Hölder continuous field l with $\langle l \rangle = 0$, there exists a $c > 0$ such that $|\langle (J \circ \varphi^n) l, \mu \rangle| \leq c \lambda^n \|l\|$, for all n . Under this assumption,

$$\left| \sum_{k=0}^{K-1} (\langle J \circ \varphi^k (a^{K-k} g + b^{K-k}), \mu \rangle - \langle J \circ \varphi^k (ag + b), \mu \rangle) \right| \leq K \|g\| c c_1 \lambda^K + K c c_2 \lambda^K, \quad (7.51)$$

because there exist constants $c_1, c_2 > 0$ such that $\|a^{K-k} - a\| \leq c_1 \lambda^{K-k}$ and $\|b^{K-k} - b\| \leq c_2 \lambda^{K-k}$ according to Lemmas 5 and 7 respectively. Letting $K \rightarrow \infty$, we conclude that the unstable contribution may be computed as

$$\langle J, \partial_s \mu_s \rangle^u = - \sum_{k=0}^{\infty} \langle J \circ \varphi^k (ag + b), \mu \rangle. \quad (7.52)$$

In practice, we need only a small number K of terms in the series above to approximate the unstable contribution to within a given tolerance. We fix some N -length μ -typical orbit $\{x_n\}$ along which we compute the first K terms above,

$$\langle J, \partial_s \mu_s \rangle^u \approx - \sum_{k=0}^{K-1} \frac{1}{N} \sum_{n=0}^{N-1} J_{n+k} (a_n g_n + b_n). \quad (7.53)$$

This computation, due to decay of correlations, converges as $N \rightarrow \infty$ and $K \rightarrow \infty$, in that order. Further, the error in approximating each term of Eq. 7.50 as an N -time ergodic average declines as $\mathcal{O}(1/\sqrt{N})$ (more details in section 9.2). The complete discussion of the error convergence of the unstable contribution computation is deferred until section 7.8.2, after the proofs of convergence of the individual computational

components of Eq. 7.53. Here we remark that even if the equivalence between Eq. 7.52 and the unstable contribution (Eq. 7.50) does not strictly hold, the error in approximating the unstable contribution (Eq. 7.50) by Eq. 7.53 becomes negligible for a sufficiently long run-up time K' (section 7.3.1). This is because, for large K' , $a^k, 0 \leq k \leq K - 1$ are all approximately equal to a and $b^k, 0 \leq k \leq K - 1$ are all approximately equal to b , due to Lemmas 5 and 7 respectively.

Now we summarize these individual components that lead to the computation of the unstable contribution by Eq. 7.53. Let $\{x_n\}_{n=-K'}^{N-1}$ be a sufficiently long finite-length μ -typical orbit. In order to compute Eq. 7.53, we need the values of a , b and g along the orbit. First, we note that the (approximate) values of a along the orbit are known from the regularized tangent solution (Eq. 7.24),

$$a_{n+1} = q_{n+1}^T ((d\varphi)_n v_n + \chi_{n+1}). \quad (7.54)$$

In section 7.5.2, we derived a recursive equation for the density gradient function,

$$g_{n+1} = \frac{g_n}{\alpha_{n+1}} - \frac{\gamma_{n+1}}{\alpha_{n+1}}, \quad (7.55)$$

which we begin by arbitrarily setting $g_{-K'} = 0$. The scalar field γ , which denotes the unstable derivative of the expansion factor α , is evaluated along an orbit by solving for the vector field w (section 7.5.1). Summarizing this step here, the recursive formula

$$w_{n+1} = \frac{1}{\alpha_{n+1}^2} (I - q_{n+1} q_{n+1}^T) ((d\varphi)_n w_n + (d^2\varphi)_n(q_n, q_n)) \quad (7.56)$$

is again started with the arbitrary choice $w_{-K'} \in \mathbb{R}^m$; at each step of the recursion, the values of γ along an orbit are set to

$$\gamma_{n+1} = \frac{q_{n+1}^T}{\alpha_{n+1}} ((d\varphi)_n w_n + (d^2\varphi)_n(q_n, q_n)). \quad (7.57)$$

Finally, the values b_n are obtained simultaneously with y_n from Eq. 7.45 and Eq.

7.49, which are repeated here for clarity:

$$y_{n+1} = \frac{(d\varphi)_n y_n + (d^2\varphi)_n(q_n, v_n)}{\alpha_{n+1}} + (d\chi)_{n+1} q_{n+1} - b_{n+1} q_{n+1} - a_{n+1} w_{n+1}, \quad (7.58)$$

where b_{n+1} is such that $y_{n+1} \cdot q_{n+1} = -v_{n+1} \cdot w_{n+1}$, at each n . Using the values of a_n , b_n and g_n obtained approximately, as outlined above, we compute an N -sample average of the k -lag correlation with J , in order to compute the k th term of the unstable contribution (Eq. 7.53). Putting this together with the stable contribution (Eq. 7.28), we obtain the overall sensitivity. These steps are condensed into the S3 algorithm listed in section 7.3.1, with one simplification: we can eliminate the need to compute Eq. 7.55 explicitly. To see this, let us consider Eq. 7.56 without the projection,

$$p_{n+1} = \frac{1}{\alpha_{n+1}^2} ((d\varphi)_n p_n + (d^2\varphi)_n(q_n, q_n)) \quad (7.59)$$

$$p_{-K'} = 0 \in \mathbb{R}^m.$$

Using these solutions p_n , Eq. 7.57 and Eq. 7.55, we find that

$$a_{n+1} w_{n+1} = a_{n+1} g_{n+1} q_{n+1} + a_{n+1} p_{n+1}.$$

Substituting this relationship into Eq. 7.58 leads to

$$y_{n+1} = \frac{(d\varphi)_n y_n + (d^2\varphi)_n(q_n, v_n)}{\alpha_{n+1}} + (d\chi)_{n+1} q_{n+1} - c_{n+1} q_{n+1} - a_{n+1} p_{n+1}, \quad (7.60)$$

where $c_{n+1} := (a_{n+1} g_{n+1} + b_{n+1})$. Notice that we only need the scalar field c , and not b and g explicitly, to compute the unstable contribution and since by iterating Eq. 7.60 we obtain c_n directly, we avoid computing g via Eq. 7.55. Since the above equation replaces the use of w_n with p_n , we do not need to solve for w_n through Eq.

7.56 either; rather, we just compute p_n by iterating Eq. 7.59. Although we used w_n in the constraint (Eq. 7.48) needed to solve for b_n , and now c_n , the constraint may be rewritten using p_n as,

$$y_{n+1} \cdot q_{n+1} = -v_{n+1} \cdot w_{n+1} = -v_{n+1} \cdot p_{n+1}. \quad (7.61)$$

The second equality is true because p_{n+1} is of the form $w_{n+1} + \delta_{n+1}q_{n+1}$, for some scalar sequence $\{\delta_{n+1}\}$, and $v_{n+1} \cdot q_{n+1} = 0$ by construction (Eq. 7.24).

7.6 Numerical results on perturbations of Baker's map

In order to validate the S3 algorithm, we consider perturbations of the standard Baker's map that are designed to elicit both stable and unstable contributions. Consider the following self-map φ_s of the torus \mathbb{T}^2 , where $s = [s_1, s_2, s_3, s_4]^T \in \mathbb{R}^4$,

$$\varphi_s([x^{(1)}, x^{(2)}]^T) = \begin{bmatrix} 2x^{(1)} + (s_1 + s_2 \sin(2x^{(2)})/2) \sin x^{(1)} - 2\pi \lfloor x^{(1)}/\pi \rfloor \\ \frac{x^{(2)}}{2} + (s_4 + s_3 \sin x^{(1)}) \sin(2x^{(2)}) + \pi \lfloor x^{(1)}/\pi \rfloor \end{bmatrix} \text{ mod } 2\pi. \quad (7.62)$$

The standard Baker's map is recovered at $s = 0 \in \mathbb{R}^4$. We isolate the effect of each parameter by illustrating the action of perturbed maps with all except one parameter set to 0. Figure 7-1 elucidates the effect of s_1 and s_2 . The unperturbed Baker's map, as shown in Figure 7-1 (top right), uniformly expands in the horizontal direction ($\hat{x}^{(1)}$) and contracts in the vertical direction ($\hat{x}^{(2)}$). By contrast, when the parameter s_1 is non-zero (and the other parameters are set to 0), the expansion in the $\hat{x}^{(1)}$ direction depends on $x^{(1)}$, resulting in a non-uniformly expanded grid as shown on the bottom left of Figure 7-1. On the bottom right of Figure 7-1, we show the effect of the parameter s_2 , from which it is clear that the expansion in the $\hat{x}^{(1)}$ direction depends nonlinearly on the $x^{(2)}$ coordinate.

In Figure 7-2 (top right), the expansion and contraction by constant factors, in the $\hat{x}^{(1)}$ and $\hat{x}^{(2)}$ directions respectively, by the unperturbed map, are clearly seen from the uniform stretching and contraction of the horizontal strips on the top left. On the bottom row, the effect of the parameters s_3 and s_4 , acting in isolation, are shown on the left and right respectively. From Eq. 7.62, it is clear that s_3 introduces a nonlinear variation with $x^{(1)}$ in the contraction along the $\hat{x}^{(2)}$ direction. This is indeed reflected in the image of the horizontal strips (top left of Figure 7-2) under the perturbed map, which is shown on the bottom left of Figure 7-2. Finally, the bottom right plot shows that the contraction in $\hat{x}^{(2)}$ is nonuniform with respect to the $x^{(2)}$ coordinate but uniform along the $x^{(1)}$ coordinate, as we would expect from Eq. 7.62 with s_4 being the only non-zero parameter.

7.6.1 Stable and unstable subspaces

The standard Baker's map has uniform stable and unstable subspaces, which are aligned with $\hat{x}^{(2)}$ and $\hat{x}^{(1)}$ respectively, i.e., $E_x^u \equiv \hat{x}^{(1)} = \text{span} \{[1, 0]^T\}$ and $E_x^s \equiv \hat{x}^{(2)} \equiv \text{span} \{[0, 1]^T\}$ at all $x \in M$. The perturbations of s_1 and s_4 do not alter these uniform stable and unstable directions. It can be verified that the perturbation of s_2 alters the stable direction while retaining $\hat{x}^{(1)}$ as the unstable direction everywhere. The perturbed map with s_3 being the only non-zero parameter has a non-uniform unstable subspace that is not parallel to $\hat{x}^{(1)}$ everywhere. On the other hand, its stable subspaces are everywhere parallel to $\hat{x}^{(2)}$.

7.6.2 SRB measures of perturbed Baker's maps

The perturbations of the Baker's map and the original map are all uniformly hyperbolic. The SRB measures of the perturbed maps associated to each parameter acting in isolation are shown in Figure 7-3. The SRB measure of the unperturbed Baker's map ($s = 0 \in \mathbb{R}^4$) is the Lebesgue measure on \mathbb{T}^2 . However, the perturbed maps have SRB measures that are not absolutely continuous with respect to Lebesgue on \mathbb{T}^2 . This can be visually observed in Figure 7-3, where, in every case, the distribution




Figure 7-1: Top left: the domain \mathbb{T}^2 covered by rectangles. The other figures show the application of Baker's map at $s = 0 \in \mathbb{R}^4$ (top right), $s = [0.2, 0, 0, 0]^T$ (bottom left) and $s = [0, 0.2, 0, 0]$ (bottom right) on the gridded top left figure.




Figure 7-2: Top left: the domain \mathbb{T}^2 covered by rectangles. The other figures show the application of Baker's map at $s = 0 \in \mathbb{R}^4$ (top right), $s = [0, 0, 0.2, 0]^T$ (bottom left) and $s = [0, 0, 0, 0.2]$ (bottom right) on the gridded top left figure.




Figure 7-3: Each plot shows the SRB distribution achieved at the parameter values indicated on the title.

appears non-smooth in the $\hat{x}^{(2)}$ direction. In fact, the SRB distribution in each case is absolutely continuous on the unstable manifold, which is uniform and parallel to $\hat{x}^{(1)}$ in every case except the s_3 perturbation. In the perturbed map with s_3 being non-zero (bottom left of Figure 7-3), the SRB measure appears smooth along its curved unstable manifold, and rough along its stable manifold, which is uniform and parallel to the $\hat{x}^{(2)}$ direction. As noted in section 7.6.1, while the unstable manifold of the perturbed map with $s_2 \neq 0$ is parallel to $\hat{x}^{(1)}$ everywhere, the stable manifold is curved around the $\hat{x}^{(2)}$ direction. This picture is consistent with smoothness of the map's SRB measure (top right of Figure 7-3) in the $\hat{x}^{(1)}$ direction and the lack of smoothness in the vertical direction, which appears distinct from the s_4 case where the stable manifold is uniformly parallel to the $\hat{x}^{(2)}$ direction.




Figure 7-4: Ergodic average of the objective function $J = \cos(4x_2)$ as a function of s_4 (left) and s_1 (right); the other parameters are set to 0. Sensitivities from S3 are shown in black at select parameter values.




Figure 7-5: Ergodic average of the objective function $J = \cos(4x_2)$ as a function of s_4 (left) and $s_1 = s_3$ (right); the other parameters are set to 0. Sensitivities from S3 are shown in black at select parameter values.

7.6.3 Sensitivities of a smooth objective function

We make an arbitrary choice of a smooth objective function, $J([x^{(1)}, x^{(2)}]^T) = \cos 4x^{(2)}$, and validate the parametric derivatives of its statistics computed by S3 in each perturbed Baker's map. The validation results are shown in Figures 7-4 and 7-5. In the left plot of Figure 7-4, only perturbations of s_1 are considered, while other parameters are held at 0. This leads to only an unstable contribution to the sensitivity because $\chi_{\varphi_{s_1} x} = [\sin x^{(1)} \ 0]^T$ which is aligned with the unstable direction $\hat{x}^{(1)} = [1 \ 0]^T$ at all x . The ergodic/ensemble averages of J are shown as a function of s_1 in blue; superimposed as black lines are the linearly extrapolated values of the S3 sensitivities computed at many different values of s_1 . We see that the derivatives of the $\langle J \rangle$ vs. s_1 are closely approximated by S3. This validates the computation of sensitivities to unstable perturbations by S3.

Similarly, we show the validation of the S3 algorithm to stable perturbations, on

the right plot of Figure 7-4. To generate these plots, ergodic/ensemble averages of J are computed at perturbed maps where only s_4 is varied. In this case, χ is aligned with $\hat{x}^{(2)}$, which is the uniform stable direction. The S3 sensitivities, shown in black, are perfectly tangent to the response curves $\langle J \rangle$ vs. s_4 , as shown in Figure 7-4 (right).

To test S3 on maps with non-uniform stable and unstable directions, we apply the algorithm to compute linear response in maps with non-zero s_2 and s_3 . At a non-zero s_2 parameter, when all other parameters are held at 0, the stable manifold is non-uniform, while the unstable manifold is uniform and tangent to $\hat{x}^{(1)}$. The overall sensitivity contains only an unstable contribution, which is verified to be correct in Figure 7-5 (left). As a final test, we put $s_1 = s_3$ and vary this parameter, which makes the unstable manifold curved and nonuniform. The response curves are shown (in blue) in Figure 7-5, on which the S3 sensitivities (black), which have both non-zero stable and unstable contributions, are plotted. Once again, the S3 derivatives are accurate over a range of parameter values. In each case, we chose $K = 11$ (terms summed in Ruelle's series) and $N = 500000$ (samples to compute ergodic averages), in the S3 algorithm; the blue points in each figure were ensemble averages over 160 million samples. The source code for this section that includes the implementation of S3 on the perturbed Baker's maps can be found at [59].

7.7 Proof of Theorem 1

In this section, we prove Theorem 2, which establishes that the S3 decomposition of Ruelle's formula into stable and unstable contributions (Eq. 7.14) is well-defined. We show that the S3 decomposition exists and is differentiable on the unstable manifold. In particular, we prove that the regularized perturbation field, v , that is the limit of the sequence $\{v^n\}$ in Eq. 7.12, exists. This proves Theorem 2-1. The second statement, Theorem 2-2., is shown by proving that this vector field v is differentiable on the unstable manifold. We then prove that the differentiability in the unstable direction of v implies that of the scalar field a , and hence Theorem 2-2. is proved. First we start by proving the existence and uniqueness of v . Before we begin, we

establish some notation for oblique projection operators on vector fields, and describe their properties that we use in the proofs.

Notation 4. Let $\mathcal{S}_x : T_x M \rightarrow T_x M$ denote the linear operator that gives the component of a tangent vector in E_x^s , in a direct sum of components along E_x^u and E_x^s . If $v_x = v_x^u + v_x^s$ where $v_x^u \in E_x^u$ and $v_x^s \in E_x^s$, then, $\mathcal{S}_x v_x := v_x^s$. Applying \mathcal{S}_x to every point on M would result in a linear operator on vector fields of M , hereafter denoted \mathcal{S} .

Remark 4. Note that \mathcal{P}_x is an orthogonal projector. Its norm is therefore 1. In contrast, the operator \mathcal{S}_x is not an orthogonal projector; \mathcal{S}_x is uniformly bounded over M , i.e., $S := \sup_{x \in M} \|\mathcal{S}_x\|$ is finite.

Remark 5. It follows from their definition that both operators are idempotent, i.e., $\mathcal{P}^2 = \mathcal{P}$ and $\mathcal{S}^2 = \mathcal{S}$.

Remark 6. Due to the covariance of E^s and E^u , i.e., $d\varphi_x^n E_x^s = E_{\varphi^n x}^s$ and $d\varphi_x^n E_x^u = E_{\varphi^n x}^u$, the operator \mathcal{S} satisfies $d\varphi_x^n \mathcal{S}_x = \mathcal{S}_{\varphi^n x} d\varphi_x^n$. Operating on vector fields on M , this equality leads to $d\varphi^n \mathcal{S} = \mathcal{S} d\varphi^n$.

Remark 7. At every $x \in M$, both \mathcal{S}_x and \mathcal{P}_x map E_x^u to 0 by their definitions. Also by definition, for any $v_x \in T_x M$, both $v_x - \mathcal{S}_x v_x$ and $v_x - \mathcal{P}_x v_x$ are in E_x^u . We thus have $\mathcal{P}_x v_x = \mathcal{P}_x \mathcal{S}_x v_x$ and $\mathcal{S}_x v_x = \mathcal{S}_x \mathcal{P}_x v_x$. The equivalent expressions, as operators on vector fields, are $\mathcal{P} \mathcal{S} = \mathcal{P}$ and $\mathcal{S} \mathcal{P} = \mathcal{S}$.

We now prove that a regularized perturbation field v exists, which is the limit of the tangent equation solved with repeated projections (section 7.4.2). We note that, in all the proofs that follow, the constants C and λ refer to the eponymous constants in the definition of uniform hyperbolicity (section 7.2), and the constant S refers to the norm of the stable projection operator \mathcal{S} , as defined above (Remark 4); all other constants, such as A, c , etc. may vary from line to line.

7.7.1 Existence and uniqueness of a regularized perturbation field

Lemma 3. *For any bounded vector field $\chi : M \rightarrow \mathbb{R}^m$, there exists a unique bounded vector field v on M that satisfies*

$$v = \mathcal{P}(d\varphi v + \chi) \quad (7.63)$$

Note that because $\mathcal{P}^2 = \mathcal{P}$, Eq. 7.63 implies that $\mathcal{P}v = v$. In other words, at every $x \in M$, v_x is orthogonal to the one-dimensional unstable subspace E_x^u . Also, a vector field χ being bounded means that $\sup_{x \in M} \|\chi_x\|$ is finite.

Proof. We first prove existence. Let $v^0 = 0$ be an arbitrary bounded vector field on M . Let $v^{k+1} = \mathcal{P}(d\varphi v^k + \chi)$ for $k = 0, 1, \dots$. We will show that for each $x \in M$, $\{v_x^k, k = 0, 1, \dots\}$ is a Cauchy sequence and thus converges to a limit, namely v_x .

From the definition of v^k and the linearity of the operators involved, we have $v^{k+1} - v^k = \mathcal{P}d\varphi(v^k - v^{k-1}) = (\mathcal{P}d\varphi)^k(v^1 - v^0)$. Here, we can use the relations between \mathcal{P} , \mathcal{S} , and $d\varphi$ (See Remarks 6 and 7) to simplify the linear operator as follows. Since $\mathcal{P} = \mathcal{P}\mathcal{S}$, $(\mathcal{P}d\varphi)^k = (\mathcal{P}\mathcal{S}d\varphi)^k = (\mathcal{P}d\varphi\mathcal{S})^k$, where the second equality follows from Remark 6. Now using $\mathcal{S}\mathcal{P} = \mathcal{S}$ (Remark 7), $(\mathcal{P}d\varphi)^k = (\mathcal{P}d\varphi\mathcal{S})^k = \mathcal{P}d\varphi(d\varphi\mathcal{S})^{k-1}$, and finally using Remark 6, $(\mathcal{P}d\varphi)^k = \mathcal{P}d\varphi(d\varphi\mathcal{S})^{k-1} = \mathcal{P}d\varphi^k\mathcal{S}$.

From inequality 7.4 and the uniform boundedness of $\|\mathcal{S}_x\|$ by S (Remark 4), $\|\mathcal{P}d\varphi^k\mathcal{S}_x\| \leq C\lambda^k S$. Together with $\|\mathcal{P}_x\| = 1$, we have a uniform bound

$$\|(\mathcal{P}d\varphi)_x^k\| = \|\mathcal{P}_{\varphi^{k_x}}d\varphi_x^k\mathcal{S}_x\| \leq C\lambda^k S, \quad \forall x \in M. \quad (7.64)$$

Thus,

$$\|v_x^{k+1} - v_x^k\| = \|(\mathcal{P}d\varphi)_{\varphi^{-k_x}}^k(v_{\varphi^{-k_x}}^1 - v_{\varphi^{-k_x}}^0)\| \leq C\lambda^k S \|v_{\varphi^{-k_x}}^1 - v_{\varphi^{-k_x}}^0\|. \quad (7.65)$$

Because v^0 is bounded by definition, and v^1 is bounded due to the boundedness of v^0, χ , and $d\varphi$, there exists an $A > 0$ such that $C S \|v_x^1 - v_x^0\| < A$ for all $x \in M$.

Thus, for all $k \geq 0$, and $x \in M$,

$$\|v_x^{k+1} - v_x^k\| \leq A \lambda^k. \quad (7.66)$$

And, for all $m > n \geq 0$,

$$\|v_x^m - v_x^n\| \leq \sum_{k=n}^{m-1} \|v_x^{k+1} - v_x^k\| \leq A \sum_{k=n}^{m-1} \lambda^k < A \frac{\lambda^n}{1-\lambda}. \quad (7.67)$$

Then, given $\epsilon > 0$, choosing $N > \lceil \log |A/(\epsilon(1-\lambda))| / \log |1/\lambda| \rceil$, for all such that $m > n \geq N$,

$$\|v_x^m - v_x^n\| \leq \epsilon, \quad \forall x \in M. \quad (7.68)$$

Thus $\{v_x^n\}$ is a *uniform* Cauchy sequence, and thus the sequence $\{v^n\}$ converges uniformly on M .

Equation 7.67, applied to $n = 0$, also implies that $\|v_x^m\| \leq \|v_x^0\| + \frac{A}{1-\lambda}$, for all x and m . The limit of the Cauchy sequence thus uniformly satisfies the same bound.

We now show this limit, defined by the vector field v as $v_x := \lim_{n \rightarrow \infty} v_x^n$, is **unique** in satisfying Eq. 7.63. Suppose Δv is a bounded vector field such that $v + \Delta v$ also satisfies Eq. 7.63. Without loss of generality, let Δv be bounded by 1, i.e., $\|\Delta v_x\| \leq 1$ for all $x \in M$. Due to the linearity of the operators involved, $\Delta v = \mathcal{P} d\varphi \Delta v$, and by iteration, $\Delta v = (\mathcal{P} d\varphi)^k \Delta v$ for any $k \in \mathbb{N}$. Using Eq. 7.64, $\|(\Delta v)_x\| \leq C \lambda^k S \|(\Delta v)_{\varphi^{-k}x}\| \leq C \lambda^k S$ for all $x \in M$. Because $\lambda < 1$ and this inequality holds for all k , $\|\Delta v_x\|$ must be 0. Thus $\Delta v = 0$. The uniqueness of v follows. \square

This establishes the existence of the regularized perturbation field v , and completes the proof of Theorem 2-1. Having shown that v exists, we now show that it is differentiable in the unstable direction. This differentiability is used to prove Theorem 2-2. Before this, we first show that q is differentiable in its own direction. Although it is known that self-derivatives of the unstable/stable subspaces exist (see, for instance,

Remark after Lemma 19.1.7 of [154]), we include a proof here for completion.

7.7.2 Existence of self-derivative of the unstable direction

We refer to as the unstable self-derivative, the vector field $w := d_\xi^2 \Phi^x$. We now show its existence.

Lemma 4. *There exists a unique, bounded vector field w that satisfies*

$$w = \mathcal{P} \frac{d^2 \varphi(q, q) + d\varphi w}{\alpha^2}, \quad (7.69)$$

where $\alpha := \|d\varphi q\|$.

Note that, since $w = \mathcal{P}w$, w is orthogonal to q .

Proof. Consider a sequence of vector fields $\{w^n\}$ that satisfies the recurrence relation

$$w^{n+1} = \mathcal{P} \frac{d^2 \varphi(q, q) + d\varphi w^n}{\alpha^2} \quad (7.70)$$

We will now show that w^n is a uniformly Cauchy sequence and hence converges uniformly. From the above recurrence relation,

$$w^{n+1} - w^n = \frac{1}{\alpha^2} \mathcal{P} d\varphi (w^n - w^{n-1}), \quad (7.71)$$

iterating which, gives, for all $n \in \mathbb{N}$

$$w^{n+1} - w^n = \frac{1}{\prod_{k=0}^{n-1} \alpha_{\varphi^{-k}}^2} (\mathcal{P} d\varphi)^n (w^1 - w^0). \quad (7.72)$$

Then, using i) $(\mathcal{P} d\varphi)^n = \mathcal{P} d\varphi^n \mathcal{S}$, which is shown in Lemma 3, and ii) $\|\mathcal{P}_x\| = 1$, at all $x \in M$,

$$\|w_x^{n+1} - w_x^n\| \leq \frac{1}{\prod_{k=0}^{n-1} \alpha_{\varphi^{-k}x}^2} \|d\varphi_{\varphi^{-n}x}^n \mathcal{S}_{\varphi^{-n}x} (w_{\varphi^{-n}x}^1 - w_{\varphi^{-n}x}^0)\|. \quad (7.73)$$

Now using i) $\prod_{k=0}^{n-1} \alpha_{\varphi^{-k}x} \geq (1/C)\lambda^{-n}$, and ii) $\|d\varphi_x^n \mathcal{S}_x\| \leq C \lambda^n S$, for all $x \in M$, both of which follow from the definition of uniform hyperbolicity,

$$\|w_x^{n+1} - w_x^n\| \leq C^3 \lambda^{3n} S \|w_{\varphi^{-n}x}^1 - w_{\varphi^{-n}x}^0\|. \quad (7.74)$$

Clearly, since the map from w^0 to w^1 is bounded, and by assumption w^0 is bounded, there exists a constant A such that $A := \sup_{x \in M} \|w_x^1 - w_x^0\|$. Hence, for any $m > n \geq 0$,

$$\|w_x^m - w_x^n\| \leq A S C^3 \sum_{k=n}^{m-1} \lambda^{3k} \leq A S C^3 \frac{\lambda^{3n}}{1 - \lambda^3}. \quad (7.75)$$

Since the above inequality holds for all x , the sequence w_x^n is uniformly Cauchy and hence converges uniformly. Let $w := \lim_{n \rightarrow \infty} w^n$. To show that w is unique, suppose w^n and \tilde{w}^n are two different sequences that both satisfy Eq. 7.70. Then, at every $x \in M$,

$$\|w_x^n - \tilde{w}_x^n\| \leq \frac{1}{\alpha_x^2} \|d\varphi_{\varphi^{-1}x} \mathcal{S}_{\varphi^{-1}x} (w_{\varphi^{-1}x}^{n-1} - \tilde{w}_{\varphi^{-1}x}^{n-1})\| \quad (7.76)$$

$$\leq A' \lambda^{3n}, \quad (7.77)$$

where the second inequality is obtained by recursively applying the first n times and then applying the definition of uniform hyperbolicity as done previously. Taking the limit $n \rightarrow \infty$ on both sides, we obtain that at every $x \in M$, $\|w_x - \tilde{w}_x\| = 0$. Thus, w is a unique vector field independent of w^0 . \square

We note that the vector field w that satisfies Eq. 7.69 is also $\partial_\xi q$. This relationship can be derived by differentiating with respect to ξ the following equation that expresses the definition of the unstable CLV, q :

$$\alpha_{\varphi x} q_{\varphi x} = d\varphi_x q_x. \quad (7.78)$$

We derive this relationship in detail in section 7.5.1. As described in section 7.5, the practical computation of the vector field w involves repeated application of Eq. 7.70,

by choosing, without loss of generality, $w_x^0 = 0$, at all x . Choosing an initial phase point x sampled according to μ , one then obtains the values $w_{\varphi^n x}^n$ along the orbit of x . As n increases, these values exponentially approach the true value of w along that orbit, as the proof of Lemma 4 shows. The convergence of the numerical procedure to compute w , which we have shown here, is needed to prove the convergence of the method (in the S3 algorithm, section 7.3.1) to compute the unstable derivative of v . We focus on this convergence next, which ultimately establishes the differentiability of the S3 decomposition.

7.7.3 Convergence of the unstable projections of the regularized tangent solutions

The S3 decomposition of the perturbation field χ into $a q$ and $\chi - a q$ is differentiable on the unstable manifolds, if $y := \partial_\xi v$ exists. To see why, we first note that the scalar field a is the limit of the iterative projections of the sequence v^n , which in practice is obtained by solving the regularized tangent equation (Eq. 7.24). We formally establish the existence of a by showing that the iterative procedure used to obtain it converges.

Lemma 5. *Let $\{a^n\}$, $n \in \mathbb{Z}^+$, be a sequence of scalar fields determined by a sequence of vector fields, $\{v^n\}$, that satisfies i) $v_x^n \cdot q_x = 0$, $\forall x \in M$, $n \in \mathbb{Z}^+$, and,*

$$ii) v^{n+1} = (d\varphi) v^n + \chi - a^{n+1} q, \quad n \in \mathbb{Z}^+. \quad (7.79)$$

Then, $\{a^n\}$ converges uniformly.

Note that i) and ii) above constitute the iterative procedure to compute the regularized tangent solution in the S3 algorithm (section 7.3.1). The sequence $\{v^n\}$ is identical to the sequence defined in Lemma 3.

Proof. The scalar fields a^{n+1} are the projections of $d\varphi v^n + \chi$ on q ,

$$a^{n+1} = q^T (d\varphi v^n + \chi). \quad (7.80)$$

As shown in Lemma 3, the sequence v^n converges uniformly. Hence, also using the fact that $\varphi \in \mathcal{C}^3(M)$, for every $\epsilon > 0$, there exists an $N \in \mathbb{N}$ such that for all $m, n \geq N$,

$$\|v_x^m - v_x^n\| < \frac{\epsilon}{\sup_{x \in M} \|d\varphi_x\|}, \quad \forall x \in M. \quad (7.81)$$

Hence, for all $m, n \geq N$, and for all $x \in M$,

$$\begin{aligned} |a_x^{m+1} - a_x^{n+1}| &\leq \|q_x\| \|d\varphi_x\| \|v_x^m - v_x^n\| \\ &\leq \left(\sup_{x \in M} \|d\varphi_x\| \right) \|v_x^m - v_x^n\| \leq \epsilon. \end{aligned} \quad (7.82)$$

Thus, $\{a^n\}$ converges uniformly. \square

We define the limit of the sequence $\{a^n\}$ by taking the limit as $n \rightarrow \infty$ of Eq. 7.80,

$$a := \lim_{n \rightarrow \infty} a^n = q^T(d\varphi v + \chi), \quad (7.83)$$

where $v := \lim_{n \rightarrow \infty} v^n$ is as defined in Lemma 3. In other words, we have the following relationship between a and v ,

$$\chi - aq = v - d\varphi v. \quad (7.84)$$

As we have shown in Lemma 3, the limit v is a unique vector field, independent of the initial condition for the iteration (Eq. 7.24 or equivalently, Eq. 7.79) provided that v^0 is bounded.

7.7.4 Route to showing existence and differentiability of the S3 decomposition

So far, we have shown that a bounded vector field v exists that is orthogonal to the unstable manifold, and the asymptotic solution of a regularized tangent equation (Eq. 7.24). We showed the existence of the stable contribution by proving that the scalar

field a , which represents the component in the unstable direction in decomposing χ , is related to v (Eq. 7.84). Further, we established that a is achieved in the same iterative procedure used for v .

In order to complete the proof of Theorem 2, we must show that the S3 decomposition of χ is differentiable on the unstable manifold. That is, we must show that the scalar field a is differentiable on the unstable manifold. Lemma 4 is a result toward this purpose. Using Eq. 7.84, we can see that if v , and hence $d\varphi v$ are differentiable with respect to ξ , and since χ is differentiable, by assumption, in all directions, we can then conclude that aq is differentiable in the unstable direction by Eq. 7.84.

Further, by Lemma 4, the derivative $w := \partial_\xi q$ exists. Thus, the scalar field a is differentiable on the unstable manifold, if v is. Therefore, it remains to show the existence of $\partial_\xi v$, in order to complete the proof of Theorem 2-1 and 2.

7.7.5 Differentiability of v in the unstable direction

Proposition 2. *The regularized perturbation field v is differentiable on the unstable manifold, i.e., $\partial_\xi v$ exists.*

We prove this proposition using the following lemma.

Lemma 6. *Let $\{\zeta^n\}$ be a uniformly exponentially converging sequence of vector fields. That is, $\{\zeta^n\}$ is uniformly converging such that there exists an $A > 0$ for which $\|\zeta_x^{n+1} - \zeta_x^n\| \leq A \lambda^n$, for all $x \in M$, and $n \in \mathbb{Z}^+$. Given a bounded vector field y^0 , the sequence $\{y^n\}$ that satisfies the following recurrence,*

$$y^{n+1} = \frac{1}{\alpha} \mathcal{P} d\varphi y^n + \zeta^{n+1}, \quad \forall n \in \mathbb{Z}^+, \quad (7.85)$$

converges uniformly to a unique vector field, $y := \lim_{n \rightarrow \infty} y^n$.

Proof. Using the recurrence relation, for all $n \in \mathbb{N}$,

$$y^{n+1} - y^n = \frac{1}{\alpha} \mathcal{P} d\varphi (y^n - y^{n-1}) + (\zeta^{n+1} - \zeta^n). \quad (7.86)$$

Applying this equation iteratively,

$$y^{n+1} - y^n = \frac{1}{\prod_{k=0}^{n-1} \alpha \circ \varphi^{-k}} (\mathcal{P} d\varphi)^n (y^1 - y^0) \quad (7.87)$$

$$+ \sum_{k=0}^{n-1} \frac{1}{\prod_{j=0}^{k-1} \alpha \circ \varphi^{-j}} (\mathcal{P} d\varphi)^k (\zeta^{n-k+1} - \zeta^{n-k}). \quad (7.88)$$

Using i) the relation $(\mathcal{P} d\varphi)^n = \mathcal{P} d\varphi^n \mathcal{S}$ – which is derived in Lemma 3 using Remarks 6 and 7 –, ii) $\|\mathcal{P}_x\| = 1$, and iii) $\prod_{k=0}^{n-1} \alpha_{\varphi^{-k}x} \geq (1/C)\lambda^{-n}$ at any $x \in M$,

$$\begin{aligned} \|y_x^{n+1} - y_x^n\| &\leq C^2 \lambda^{2n} S \|y_{\varphi^{-n}x}^1 - y_{\varphi^{-n}x}^0\| \\ &\quad + C^2 S \sum_{k=0}^{n-1} \lambda^{2k} \|\zeta_{\varphi^{-k}x}^{n-k+1} - \zeta_{\varphi^{-k}x}^{n-k}\|. \end{aligned} \quad (7.89)$$

Now since $\{\zeta^n\}$ is uniformly exponentially converging,

$$\|\zeta_{\varphi^{-k}x}^{n-k+1} - \zeta_{\varphi^{-k}x}^{n-k}\| \leq A \lambda^{n-k}. \quad (7.90)$$

Further, since y^0, ζ^1 are bounded vector fields, and $\mathcal{P} d\varphi$ is bounded, y^1 is a bounded vector field. Hence, there exists some constant A_1 such that for any $x \in M$,

$$\|y_{\varphi^{-n}x}^1 - y_{\varphi^{-n}x}^0\| < A_1. \quad (7.91)$$

Using both the above relationships (7.90 and 7.91) in 7.89,

$$\|y_x^{n+1} - y_x^n\| \leq C^2 \lambda^{2n} S A_1 + A C^2 S \lambda^n / (1 - \lambda) \leq A_2 \lambda^n.$$

Hence, for $m \geq n > 0$, and all $x \in M$,

$$\|y_x^m - y_x^n\| \leq \sum_{k=n}^{m-1} A_2 \lambda^k < A_2 \frac{\lambda^n}{1 - \lambda}. \quad (7.92)$$

Thus, $\{y^n\}$ is uniformly Cauchy and converges uniformly.

To see that the limit $y := \lim_{n \rightarrow \infty} y^n$ is unique, let y^n and \tilde{y}^n be two different

bounded sequences that satisfy Eq. 7.85, and converge to y and \tilde{y} respectively. Then,

$$y^n - \tilde{y}^n = \frac{1}{\alpha} \mathcal{P} d\varphi (y^{n-1} - \tilde{y}^{n-1}), \quad (7.93)$$

which can be applied recursively to yield,

$$y^n - \tilde{y}^n = \frac{1}{\prod_{k=0}^{n-1} \alpha \circ \varphi^{-k}} (\mathcal{P} d\varphi)^n (y^0 - \tilde{y}^0) = \frac{1}{\prod_{k=0}^{n-1} \alpha \circ \varphi^{-k}} \mathcal{P} d\varphi^n \mathcal{S}(y^0 - \tilde{y}^0). \quad (7.94)$$

Since both y^0 and \tilde{y}^0 are bounded, the operator norm of \mathcal{S}_x is uniformly bounded above, and from the definition of uniform hyperbolicity, there exists some $c > 0$ such that for all n , and all $x \in M$, we have

$$\|\tilde{y}_x^n - y_x^n\| \leq c \lambda^{2n}. \quad (7.95)$$

Thus, $\lim_{n \rightarrow \infty} (\tilde{y}_x^n - y_x^n) = 0 \in \mathbb{R}^m$ for every $x \in M$. Hence, the limit y is unique. \square
 \square

Now we prove Proposition 2. We set

$$\zeta^{n+1} = \mathcal{P} \left(\frac{d^2\varphi(q, v^n)}{\alpha} + d\chi q \right) + (\partial_\xi \mathcal{P})(d\varphi v^n + \chi), \quad (7.96)$$

where the sequence $\{v^n\}$ is as defined in Lemma 3; $d^2\varphi_x(\cdot, \cdot) : T_x M \times T_x M \rightarrow T_{\varphi x} M$ is a bilinear form representing the second-order derivative of φ at $x \in M$. In order to apply Lemma 6, we must show that $\{\zeta^n\}$ is a bounded, uniformly exponentially converging sequence. First we note that the projection operator \mathcal{P} is differentiable in the unstable direction. Using the matrix representation of \mathcal{P} , and Lemma 4, its derivative is given by

$$\partial_\xi \mathcal{P} = -\partial_\xi (qq^T) = -(wq^T + qw^T), \quad (7.97)$$

from which we see that the operator norm of $\partial_\xi \mathcal{P}$ is bounded on M . Because $d\varphi \in C^2(M)$ by assumption, $\|d\varphi_x\|$ and $\|d^2\varphi_x\|$ are also bounded on M ; χ and $d\chi$ are

bounded by assumption. Now, the sequence $\{v^n\}$ is uniformly bounded on M and \mathbb{Z}^+ ($\because \{v^n\}$ is bounded and uniformly convergent, by Lemma 3). Thus, $\{\zeta^n\}$ defined in Eq. 7.96 is a uniformly bounded sequence.

Now to show that $\{\zeta^n\}$ is uniformly exponentially converging, we use its definition, Eq. 7.96, to get

$$\zeta^{n+1} - \zeta^n = \partial_\xi \mathcal{P} d\varphi (v^n - v^{n-1}) + \frac{\mathcal{P}}{\alpha} d^2\varphi (q, (v^n - v^{n-1})). \quad (7.98)$$

Since $\partial_\xi \mathcal{P}$, $d\varphi$ and $d^2\varphi$ are bounded and $\|\mathcal{P}\| = 1$, and using Eq. 7.65, there exists some $c > 0$ such that, for all $x \in M$,

$$\|\zeta_x^{n+1} - \zeta_x^n\| \leq \left(\|(\partial_\xi \mathcal{P})_x\| \|d\varphi_{\varphi^{-1}x}\| + \frac{\|d^2\varphi_x(q_x, \cdot)\|}{\alpha} \right) A \lambda^n := c \lambda^n. \quad (7.99)$$

Thus, using the same argument following Eq. 7.65 in Lemma 3, we can show that $\{\zeta^n\}$ converges uniformly. Further, the limit $\zeta := \lim_{n \rightarrow \infty} \zeta^n$ is unique, and defined by taking the limit of Eq. 7.96,

$$\zeta = \mathcal{P} \left(\frac{d^2\varphi(q, v)}{\alpha} + d\chi \cdot q \right) + (\partial_\xi \mathcal{P})(d\varphi v + \chi). \quad (7.100)$$

Therefore, $\{\zeta^n\}$, as defined in Eq. 7.96, is a uniformly bounded, uniformly exponentially converging sequence of vector fields. Now applying Lemma 6, there exists a unique, bounded vector field y that satisfies

$$y = \frac{\mathcal{P} d\varphi y}{\alpha} + \zeta. \quad (7.101)$$

Differentiating with respect to ξ the constraint, Eq. 7.63, which is satisfied by the regularized perturbation field v , we see that if the derivative $\partial_\xi v$ exists, it must satisfy Eq. 7.101 (replacing y), with ζ defined in Eq. 7.100. Since a vector field y that satisfies Eq. 7.101 exists and is unique (Lemma 6), the regularized perturbation field v is indeed differentiable. This completes the proof of 2. Further, the computation of its derivative, using the sequence of vector fields $\{y^n\}$ used in the proof of Lemma

6, converges exponentially along almost every trajectory. Note that the evaluation of y along a trajectory, with ζ defined in Eq. 7.100, is precisely that derived in section 7.5.3.

Recall the argument (section 7.7.4) that the differentiability (with respect to ξ) of v implies that of a . This completes the proof of Theorem 2-2., and hence the first two parts of Theorem 2. To complete the proof of Theorem 2-3, we provide a constructive proof of the unstable derivative $b = \partial_\xi a$. In a similar vein to Lemmas 3, 4 and 6, this proof shows that a trajectory-based computation adopted in S3 converges exponentially, in this case, to the true values of b .

7.7.6 Convergence of the derivative of $\{a^n\}$

We now show that the iterative algorithm for the scalar derivative $b^n := \partial_\xi a^n$ converges uniformly. Expanding Eq. 7.96 to make the appearance of $b^n := \partial_\xi a^n$ explicit, by using the definition of a^n (Lemma 5),

$$y^{n+1} = \frac{d\varphi y^n}{\alpha} + \frac{d^2\varphi(v^n, q)}{\alpha} + d\chi q - a^{n+1}w - b^{n+1}q.$$

Taking inner product with q and using $w_x \cdot q_x = 0$ at all $x \in M$ (Lemma 4),

$$b^{n+1} = \left(\frac{d\varphi y^n}{\alpha} - y^{n+1} + \frac{d^2\varphi(v^n, q)}{\alpha} + d\chi q \right) \cdot q. \quad (7.102)$$

Lemma 7. *The sequence $\{b^n\}$ defined in Eq. 7.102 converges uniformly where i) the sequence $\{y^n\}$ satisfies Eq. 7.85, and, ii) $\{v^n\}$ satisfies Eq. 7.25.*

Proof. We prove that the sequence of vector fields $\{\tau^n\}$ given by

$$\tau^{n+1} = \frac{d\varphi y^n}{\alpha} - y^{n+1} + \frac{d^2\varphi(v^n, q)}{\alpha} + d\chi q \quad (7.103)$$

converges uniformly. From the above definition, and using the linearity of $d^2\varphi(\cdot, q)$,

at any $x \in M$, and all $n > m \geq 0$,

$$\|\tau_x^{n+1} - \tau_x^{m+1}\| \leq \frac{\|d\varphi_x\|}{\alpha_x} \|y_x^n - y_x^m\| + \|y_x^{n+1} - y_x^{m+1}\| + \frac{\|d^2\varphi_x(v_x^n - v_x^m, q_x)\|}{\alpha_x}. \quad (7.104)$$

Recall that $\varphi \in C^3(M)$ and hence $\|d\varphi_x\|$ and $\|d^2\varphi_x(\cdot, \cdot)\|$ are both uniformly bounded; from uniform hyperbolicity, $1/\alpha_x \leq C\lambda$, at all $x \in M$. Let $\epsilon > 0$ be given. Using Lemma 6, there exists an $N_y \in \mathbb{N}$ such that for all $x \in M$ and $m, n \geq N_y$, $\|y_x^n - y_x^m\| < \epsilon/(3C\lambda \sup_{x \in M} \|d\varphi_x\|)$. Similarly, using Lemma 3, there exists an $N_v \in \mathbb{N}$ such that for all $x \in M$ and $m, n \geq N_v$, $\|v_x^n - v_x^m\| < \epsilon/(3C\lambda \sup_{x \in M} \|d^2\varphi_x(q, \cdot)\|)$. Choosing $N = \max\{N_v, N_y\}$, for all $m, n \geq N$, and all $x \in M$, $\|\tau_x^n - \tau_x^m\| \leq \epsilon$. Thus, $\{\tau^n\}$ converges uniformly. Since at all $x \in M$, $b_x^n = \tau_x^n \cdot q_x$, the sequence $\{b^n\}$ also converges uniformly. Moreover, note from 7.104 that the speed of convergence is exponential, since both $\{v^n\}$ and $\{y^n\}$ are exponentially uniformly converging as shown in Lemmas 3 and 6 respectively. The limit is unique since,

$$b := \lim_{n \rightarrow \infty} b^n = \left(\lim_{n \rightarrow \infty} \tau^n \right) \cdot q = \left(d\chi q - y + \frac{d\varphi y + d^2\varphi(v, q)}{\alpha} \right) \cdot q. \quad (7.105)$$

□

7.8 Proof of Theorem 2

In this section, we complete the proof of Theorem 2, which establishes the convergence of the S3 algorithm (section 7.3.1). In particular, we prove that the ergodic averages we compute in the S3 algorithm (section 7.3.1) for the stable and unstable contributions converge to their true values.

7.8.1 Convergence of the stable contribution

Having shown the S3 decomposition is differentiable on the unstable manifold (section 7.7), we now derive an alternative expression for the stable contribution. This

expression leads to a direct computation of the stable contribution using the regularized tangent equation (section 7.4.2). We then show that this computation converges to the true stable contribution.

Proposition 3. *The error in the stable contribution, when computed as an ergodic average on a trajectory of length N , converges as $\mathcal{O}(\sqrt{\log \log N}/\sqrt{N})$.*

Proof. Taking $\lim_{m \rightarrow \infty}$ of Eq. 7.67, $\|v_x^n - v_x\| < A \frac{\lambda^n}{1 - \lambda}$, $\forall x \in M$, $n \in \mathbb{Z}^+$. Hence, $\{v^n\}$ is a uniformly bounded sequence. Since dJ is bounded by assumption, by Lebesgue dominated convergence on the sequence of functions $\{dJ \cdot v^n\}$, $\lim_{n \rightarrow \infty} \langle dJ \cdot v^n, \mu \rangle = \langle dJ \cdot v, \mu \rangle$. From Lemma 3, an explicit expression for v^n starting from $v^0 = 0 \in \mathbb{R}^m$ is $v^n = \sum_{k=0}^{n-1} d\varphi^k(\chi - a^{n-k}q)$. Thus, $\langle J, \partial_s \mu_s \rangle^s = \lim_{n \rightarrow \infty} \langle dJ \cdot v^n, \mu \rangle = \langle dJ \cdot v, \mu \rangle$. Choosing a constant $c \geq \sup_{x \in M} \|(dJ)_x\|A/(1 - \lambda)$,

$$|(dJ)_x \cdot v_x^n - (dJ)_x \cdot v_x| < c \lambda^n \quad \forall x \in M, n \in \mathbb{Z}^+. \quad (7.106)$$

Let x be an arbitrary point on M chosen μ -a.e. As usual, we use the shorthand f_n to represent f_{x_n} , the vector field f evaluated at the point x_n . The error in the N -time ergodic average is as follows,

$$\left| \frac{1}{N} \sum_{n=0}^{N-1} (dJ)_n \cdot (v_n^n - v_n) \right| \leq \frac{c}{N} \sum_{n=0}^{N-1} \lambda^n := \frac{c_1}{N}. \quad (7.107)$$

Thus, the error in the stable contribution computation can be bounded as follows.

$$\begin{aligned} \left| \frac{1}{N} \sum_{n=0}^{N-1} (dJ)_n \cdot v_n^n - \langle dJ \cdot v, \mu \rangle \right| &\leq \left| \frac{1}{N} \sum_{n=0}^{N-1} (dJ)_n \cdot (v_n^n - v_n) \right| + \left| \frac{1}{N} \sum_{n=0}^{N-1} (dJ)_n \cdot v_n - \langle dJ \cdot v, \mu \rangle \right| \\ &\leq \frac{c_1}{N} + \frac{c_2 \sqrt{\log \log N}}{\sqrt{N}}. \end{aligned} \quad (7.108)$$

The bound $c_2 \sqrt{\log \log N}/\sqrt{N}$ appears since $(dJ) \cdot v$ is Hölder continuous, and the law of iterated logarithm applies to the convergence of its ergodic average [74]. Thus, we conclude that the error in the stable contribution computation converges, with N – the trajectory length or the number of sample points according to μ – almost as a

typical Monte Carlo integration,

$$\left| \frac{1}{N} \sum_{n=0}^{N-1} (dJ)_n \cdot v_n^n - \langle dJ \cdot v, \mu_s \rangle \right| \sim \mathcal{O}(\sqrt{\log \log N} / \sqrt{N}). \quad (7.109)$$

□

7.8.2 Convergence of the unstable contribution

We establish the convergence of the unstable contribution with the help of two lemmas. The first shows the convergence of the recursive formula for the logarithmic density gradient, g . The second lemma, Lemma 7, proves the convergence of the iterative procedure to differentiate the regularized perturbation field, v , on the unstable manifold. Together, these two results can be used to show that the unstable contribution, with the computation described in the S3 algorithm (section 7.3.1), converges.

Convergence of logarithmic density gradient

Lemma 8. *Given a bounded scalar field r , there exists a unique, bounded scalar field g that satisfies*

$$g \circ \varphi = \frac{g}{\alpha \circ \varphi} + r \circ \varphi, \quad (7.110)$$

Note that when $r = -\gamma/\alpha$, the logarithmic density gradient function satisfies Eq. 7.110.

Proof. First we show existence. Let $h^0 : M \rightarrow \mathbb{R}$ be an arbitrary bounded scalar function. Consider the sequence of scalar functions $\{h^n\}$ which follow the recurrence relation:

$$h^{n+1} \circ \varphi = \frac{h^n}{\alpha \circ \varphi} + r \circ \varphi. \quad (7.111)$$

We show that this sequence converges uniformly. From the above recurrence relation, at every $x \in M$,

$$h_x^{n+1} - h_x^n = \frac{h_{\varphi^{-1}x}^n}{\alpha_x} - \frac{h_{\varphi^{-1}x}^{n-1}}{\alpha_x}. \quad (7.112)$$

Applying this relation recursively, for all $n \geq 1$,

$$|h_x^{n+1} - h_x^n| \leq \frac{1}{\prod_{k=0}^{n-1} \alpha_{\varphi^{-k}x}} |h_{\varphi^{-n}x}^1 - h_{\varphi^{-n}x}^0|. \quad (7.113)$$

Under the assumption of uniform hyperbolicity, there exist constants $C > 0$ and $\lambda \in (0, 1)$ such that $\prod_{k=0}^{n-1} \alpha_{\varphi^{-k}x} \geq (1/C)\lambda^{-n}$, for all $x \in M$. Thus,

$$|h_x^{n+1} - h_x^n| \leq C\lambda^n |h_{\varphi^{-n}x}^1 - h_{\varphi^{-n}x}^0|. \quad (7.114)$$

Since h^0 is a bounded function, and h^1 is also a bounded function because h^0 and r are bounded, there exists $A > 0$ such that $|h_x^1 - h_x^0| \leq A$, for all $x \in M$. Thus, for all $i > j \geq 0$,

$$|h_x^i - h_x^j| \leq \sum_{n=j}^{i-1} |h_x^{n+1} - h_x^n| \leq AC \frac{\lambda^j}{1-\lambda}. \quad (7.115)$$

Hence $\{h^n\}$ is a uniformly Cauchy sequence and therefore converges uniformly. Let $g := \lim_{n \rightarrow \infty} h^n$. We show that this limit is unique. Let $\{h^n\}$ and $\{\tilde{h}^n\}$ be two different sequences satisfying Eq. 7.111. Additionally, assume that both h^0 and \tilde{h}^0 are bounded functions. Then,

$$|h_x^n - \tilde{h}_x^n| \leq \frac{|h_{\varphi^{-1}x}^{n-1} - \tilde{h}_{\varphi^{-1}x}^{n-1}|}{\alpha_{\varphi^{-1}x}}, \quad (7.116)$$

which by iteration, gives,

$$|h_x^n - \tilde{h}_x^n| \leq \frac{|h_{\varphi^{-n}x}^0 - \tilde{h}_{\varphi^{-n}x}^0|}{\prod_{k=1}^n \alpha_{\varphi^{-k}x}} \leq C\lambda^n |h_{\varphi^{-n}x}^0 - \tilde{h}_{\varphi^{-n}x}^0|. \quad (7.117)$$

Thus $|h_x^n - \tilde{h}_x^n| \rightarrow 0$ as $n \rightarrow \infty$ since $|h_{\varphi^{-n}x}^0 - \tilde{h}_{\varphi^{-n}x}^0| < \text{const}$. Since this holds for all $x \in M$, $\lim_{n \rightarrow \infty} h^n = \lim_{n \rightarrow \infty} \tilde{h}^n = g$. \square \square

We have shown that the limit g is independent of the initial condition h^0 . Hence, without loss of generality, we may assume that $h_x^0 = 0$ at all $x \in M$, as we do in section 7.5.2. We remark that, due to the algebraic simplification introduced in section 7.5.4, the explicit iteration of Eq. 7.111 is subsumed under the combination of Eq. 7.59 and 7.60.

Convergence of the unstable contribution computation

In the S3 algorithm, the computation of the unstable contribution is carried out as an ergodic average. We use the results we have set up so far to show that this ergodic average converges to the true unstable contribution. From our derivation (section 7.4.4) and the strong decay of correlations assumption (section 7.5.4), we obtain the following regularized expression for the unstable contribution,

$$\langle J, \partial_s \mu_s \rangle^u = - \sum_{k=0}^{\infty} \langle J \circ \varphi^k(a g + b), \mu \rangle. \quad (7.118)$$

In order to show that our computation of the unstable contribution converges, we first note that, due to exponential decay of correlations between Hölder continuous functions (section 7.4.4), the above series converges. Further, we show in Lemmas 5, 7 and 8 that the convergence to the true values of a , b and g respectively is exponential in each case, along any μ -typical orbit. Thus, iterating the equations for these quantities for a sufficient run-up time, we may assume that they are, up to machine precision, equal to their true values. The convergence rate of the unstable contribution is thus determined by that of the following N -sample averages, at the first few $k \in \mathbb{Z}^+$. For almost every x , and all $N > N_0$,

$$\left| \sum_{k=0}^{K-1} \left(\langle J \circ \varphi^k(a g + b), \mu \rangle - \frac{1}{N} \sum_{n=0}^{N-1} J_{n+k}(a_n g_n + b_n) \right) \right| \leq \frac{\sqrt{\log \log N} \sum_{k < K} c_k}{\sqrt{N}}. \quad (7.119)$$

The bound above is again obtained by applying the law of iterated logarithm, which holds for Hölder continuous functions, to the ergodic average at each k . Given $\epsilon > 0$, there exists a K such that $|\sum_{k>K} \langle J \circ \varphi^k(a g + b), \mu \rangle| < \epsilon/2$, due to the convergence of 7.118. Then, choosing N large enough such that $\sqrt{\log \log N} \sum_{k<K} c_k / \sqrt{N} < \epsilon/2$, we approximate the unstable contribution arbitrarily well,

$$\left| \langle J, \partial_s \mu_s \rangle^u - \sum_{k=0}^{K-1} (1/N) \sum_{n=0}^{N-1} J_{n+k} (a_n g_n + b_n) \right| < \epsilon.$$

This numerical approximation converges to the unstable contribution as $N \rightarrow \infty$ followed by as $K \rightarrow \infty$. We have already shown that the stable contribution computation has an error convergence rate of $\mathcal{O}(\sqrt{\log \log N} / \sqrt{N})$ (Lemma 3). This proves that the sum of stable and unstable contribution converges as $\mathcal{O}(\sqrt{\log \log N} / \sqrt{N})$, with the S3 algorithm implemented on an N -length trajectory.

7.9 Discussion and conclusion

The main contribution of this chapter is a new algorithm, called space-split sensitivity (S3), to compute linear response, or the sensitivities of statistics to parameters, in chaotic dynamical systems. The chapter presents the derivation of S3 with Ruelle's linear response formula as the starting point. While Ruelle's formula leads to an ill-conditioned direct computation, we derive a decomposition – the S3 decomposition – of the formula that yields a well-conditioned ergodic-averaging computation. The S3 decomposition is achieved by orthogonalizing the conventional tangent solution with respect to the unstable subspaces along a long trajectory. This procedure yields a regularized tangent vector field v and a scalar field a .

The sensitivity to the unstable vector field so achieved – $a q$ – is called the unstable contribution. The stable contribution is the remaining linear response, and can be calculated just as tangent sensitivities are in non-chaotic systems, but using the regularized tangent solutions in place of the conventional tangent solutions. The unstable contribution is integrated by parts on the unstable manifold to yield an expression

in which, unlike the original form of Ruelle’s formula, the integrand is bounded at all times. Then, exploiting the decay of correlations, this regularized formula can be computed fast as a sum of Monte Carlo integrals. However, the problem is, we obtain two unknown functions – the derivative of a on the unstable manifold, and the density gradient. The density gradient is a fundamental object whose regularity we link, in a different study [248], to the validity of linear response [277][276].

The density gradient is the unstable derivative of the logarithmic conditional density of the SRB measure on the unstable manifold. We take the unstable derivative of Pesin’s formula to obtain a recursive evaluation of the density gradient along a typical trajectory. We show that starting this recursion with an arbitrary scalar field leads to an exponential convergence to the true density gradient. A recursive procedure is also derived for the unstable derivative of a by combining the unstable derivatives of v and q . The evaluation of the latter two unstable derivatives requires solving two second-order tangent equations, which are the most computationally intensive steps of the S3 algorithm.

Overall, the S3 algorithm computes Ruelle’s formula as a well-conditioned ergodic average. The efficiency of the S3 algorithm stems from its decomposition of Ruelle’s formula and its treatment of the unstable contribution. In Ruelle’s formula (Eq. 7.11), the norm of the integrand increases exponentially with n , which leads to the ill-conditioning of the estimation of each term as an ergodic average. However, in the unstable contribution after applying the S3 decomposition (Eq. 7.34), the norm of the integrand does not grow with n and hence, the estimation of the S3 formula as an ergodic average is efficient.

An important direction of future work is to make S3 applicable to systems with higher-dimensional unstable manifolds.

Chapter 8

Extending S3 to Chaotic ODEs

In this chapter, we extend the S3 algorithm for linear response, developed in the previous chapter for discrete-time systems, to ODEs. We consider chaotic ODEs which possess attractors that have a one-dimensional unstable manifold. In comparison to uniformly hyperbolic maps, which we studied in the previous chapter, the continuous-time systems we consider here, also known as uniformly hyperbolic flows, are complicated by the presence of a one-dimensional center subspace of the tangent space, which is neither contracting nor expanding exponentially under the tangent dynamics. We split Ruelle's formula into three parts: the stable, unstable and center contributions. The stable and unstable contributions have similar forms as in the case of uniformly hyperbolic maps, although they are different in the present setting of uniformly hyperbolic flows. We derive expressions and outline the computation of the three components of linear response. An ergodic-averaging scheme, once again with an error convergence at the rate of a typical Monte Carlo summation of Ruelle's response, is presented. All the notations and definitions from Chapter 7 carry over.

8.1 Chaotic ODE

In this chapter, we consider a time-discretization of a chaotic governing equation of the form

$$\frac{dx}{dt} = F(x; s). \quad (8.1)$$

The ODE, generated by the vector field F , is numerically integrated in time. As usual, we use φ_s to denote integration by one timestep, which is assumed to be small enough such that $\varphi_s : M \rightarrow M$ is a C^3 function. Let f be the unit vector field in the tangent subspace spanned by F i.e., $f_x = F_x / \|F_x\|$, at all $x \in M$. We also introduce the scalar field β to represent the norm of the generator of the flow, $\beta_x := \|F_x\|$.

8.2 Uniformly hyperbolic flows

The purpose of this chapter is to extend the S3 algorithm, which was presented in chapter 7 for uniformly hyperbolic systems, to uniformly hyperbolic flows, in which linear response also holds [235][85].

In the class of systems we consider, the tangent space at every $x \in M$ can be decomposed into three $d\varphi$ -invariant or covariant subspaces, namely, i) E_x^u – the unstable subspace, ii) E_x^s – stable subspace and iii) $E_x^c \equiv \text{span}\{f\}$ – center subspace. Since both E_x^u and E_x^c are one-dimensional at all x , this means that, for any $v \in T_x M$, a direct sum decomposition of the form, $v = a^u q + v^s + a^c f$ exists, where $v^s \in E_x^s$ and a^u and a^c are scalar components in the unstable and center directions respectively.

8.2.1 The center subspace

Considering a small fixed timestep, δt , the solution map for Eq. 2.1 is approximately $\varphi_s : M \rightarrow M$ given by,

$$\varphi_s x = x + \delta t F(x, s). \quad (8.2)$$

As usual, we write φ_{s_0} simply as φ . The tangent equation, corresponding to an infinitesimal perturbation of s ,

$$v_{n+1} = d\varphi v_n + \chi_{n+1}, \quad (8.3)$$

is in the same form as in uniformly hyperbolic diffeomorphisms, where $\chi_{n+1} = (\delta t)\partial_s F(x_n, s)$. The stable and unstable subspaces consist of uniformly exponentially decaying tangent vectors, under $d\varphi^n$ and $d\varphi^{-n}$ respectively, as described in section 6.2.2. The one-dimensional center subspace is spanned by the right-hand side of the ODE (2.1), F . It is covariant, i.e., $d\varphi_x E_x^c = E_{\varphi x}^c$. Further, it does not exhibit growth or decay on the exponential scale, under the tangent dynamics. Using the covariance of E_x^c and the definition of β , $\|d\varphi_x f_x\| = \|(\beta_{\varphi x}/\beta_x) f_{\varphi x}\|$. Since $F \in C^3(M)$ by assumption, $\|\beta\| := \sup_{x \in M} \|F_x\|$ is finite. We assume that $\|\beta\|_* := \inf_{x \in M} \|F_x\| > 0$. Then, $\|d\varphi|_{E^c}\| := \sup_{x \in M} \|d\varphi_x f_x\| \leq \|\beta\| / \|\beta\|_* := B$, is finite.

8.3 Center-Unstable projection

We define the operator \mathcal{P} on vector fields to project out of the center-unstable subbundle. That is, for any $v \in E^u \oplus E^c$, $\mathcal{P}v = 0$. We describe some properties of \mathcal{P} .

Remark 8. *Considering the decomposition $v = a^u q + a^f f + \mathcal{S}v$, where a^u and a^f are scalar field components of v along q and f respectively, $\mathcal{P}v = \mathcal{P}\mathcal{S}v$ since $\mathcal{P}q = \mathcal{P}f = 0$. As operators, $\mathcal{P} = \mathcal{P}\mathcal{S}$.*

Remark 9. *For any vector field v , $(I - \mathcal{P})v \in E^u \oplus E^c$. Thus, $\mathcal{S}(I - \mathcal{P}) = 0$, or $\mathcal{S} = \mathcal{S}\mathcal{P}$.*

Remark 10. *Using the above remarks, we can simplify the operator $(\mathcal{P}d\varphi)^n$ as follows. Since $\mathcal{P} = \mathcal{P}\mathcal{S}$ (Remark 8), $(\mathcal{P}d\varphi)^n = \mathcal{P}\mathcal{S}d\varphi(\mathcal{P}d\varphi)^{n-1}$, and then using the covariance of \mathcal{S} , $(\mathcal{P}d\varphi)^n = \mathcal{P}d\varphi\mathcal{S}(\mathcal{P}d\varphi)^{n-1}$. Using $\mathcal{S}\mathcal{P} = \mathcal{S}$ (Remark 9), and the covariance of \mathcal{S} , $(\mathcal{P}d\varphi)^n = \mathcal{P}d\varphi^2\mathcal{S}(\mathcal{P}d\varphi)^{n-2}$. Repeatedly applying $\mathcal{S}\mathcal{P} = \mathcal{S}$ and the covariance of \mathcal{S} , $n - 2$ times, we obtain, $(\mathcal{P}d\varphi)^n = \mathcal{P}d\varphi^n\mathcal{S}$.*

First we prove the existence of a regularized tangent solution, analogous to the uniformly hyperbolic case (Lemma 3).

Lemma 9. *For any bounded vector field χ , there exists a unique, bounded vector field v that satisfies*

$$v = \mathcal{P}(d\varphi v + \chi). \quad (8.4)$$

The proof, carried out using Remark 10, is identical to that of Lemma 3 for uniformly hyperbolic maps without the center direction. The speed of uniform convergence of $\{v^n\}$ (Lemma 3) is also exponential. We convert the above lemma into a numerical procedure for evaluating v along a trajectory. As before, we refer to v as the regularized tangent solution, and we compute it by regularizing the conventional tangent solution. Recall that, fixing a μ -typical trajectory, in section 7.4.3, we derived the following regularized tangent equation,

$$v_{n+1} = (d\varphi)_n v_n + \chi_{n+1} - a_{n+1} q_{n+1}, \quad (8.5)$$

where a_{n+1} is chosen as to make $v_{n+1} \cdot q_{n+1} = 0$. We now project out the center direction as well by adding a term $c_{n+1} f_{n+1}$ so that

$$v_{n+1} = (d\varphi)_n v_n + \chi_{n+1} - a_{n+1} q_{n+1} - c_{n+1} f_{n+1}, \quad (8.6)$$

where a_{n+1} and c_{n+1} are chosen so as to make $v_{n+1} \cdot f_{n+1} = 0$, and simultaneously, $v_{n+1} \cdot q_{n+1} = 0$. Explicitly, in field notation,

$$c = \frac{v^+ \cdot (f - \theta q)}{1 - \theta^2} \quad (8.7)$$

$$a = \frac{v^+ \cdot (q - \theta f)}{1 - \theta^2}, \quad (8.8)$$

where $v^+ := d\varphi v + \chi$ and $\theta := q \cdot f$. Without loss of generality, we can consider $v^0 = 0$

in Lemma 9. Then, we can write v^n down in closed form as,

$$v^n = \sum_{k=0}^{n-1} d\varphi^k (\chi - a^{n-k}q - c^{n-k}f) = \sum_{k=0}^{n-1} \mathcal{P}d\varphi^k \chi. \quad (8.9)$$

Due to the exponential convergence of v^n , if the run-up time is sufficiently long, the above regularized tangent equation (Eq. 8.6) evolves arbitrarily close to the true values of v , a and c along a trajectory.

8.4 Stable–Unstable–Center decomposition

We propose the following decomposition of Ruelle’s formula, motivated by the regularized tangent solution (Eq. 8.9)

$$\begin{aligned} \langle J, \partial_s \mu_s \rangle &= \lim_{K \rightarrow \infty} \sum_{k=0}^{K-1} \langle d(J \circ \varphi^k) \cdot (a^{K-k}q + c^{K-k}f), \mu \rangle \\ &+ \lim_{K \rightarrow \infty} \sum_{k=0}^{K-1} \langle d(J \circ \varphi^k) \cdot (\chi - a^{K-k}q - c^{K-k}f), \mu \rangle. \end{aligned} \quad (8.10)$$

Here the first term on the right hand side is the unstable-center contribution, while the second is the stable contribution. Analogous to the proof in the previous chapter, we can show that the stable contribution can be alternatively expressed as,

$$\lim_{K \rightarrow \infty} \sum_{k=0}^{K-1} \langle d(J \circ \varphi^k) \cdot (\chi - a^{K-k}q - c^{K-k}f), \mu \rangle = \langle dJ \cdot v, \mu \rangle, \quad (8.11)$$

where v is orthogonal to $E^u \oplus \text{span}\{f\}$. The computation left to be tackled is that of the center-unstable contribution, which we split further into the unstable contribution and the center contribution. The center contribution, which we turn our attention to first, is defined as,

$$\langle J, \partial_s \mu_s \rangle^c = \lim_{K \rightarrow \infty} \sum_{k=0}^{K-1} \langle d(J \circ \varphi^k) \cdot (c^{K-k}f), \mu_s \rangle. \quad (8.12)$$

In the above expression, the integrand, $d(J \circ \varphi^k) \cdot f$, may increase in norm, albeit sub-exponentially. Thus, the variance of evaluating the ensemble average as a finite-sample ergodic average of tangent solutions increases with k , since numerically, small errors will increase in the unstable direction. Hence, we resort to evaluating an equivalent expression for the center contribution, without time-evolving tangent solutions. In particular, the ergodic average of the equivalent expression must not have a variance that increases with k .

8.5 Center contribution alternative expression

We now derive an alternative expression for the center contribution that is easier to compute than Eq. 8.12. We may assume that iterating the regularized tangent equation for a sufficiently long run-up time ensures that the scalar field c has converged. Thus, we may consider the center contribution to be equivalent to the following expression, without the superscript on c ,

$$\langle J, \partial_s \mu_s \rangle^c = \sum_{n=0}^{\infty} \langle d(J \circ \varphi^n) \cdot (cf), \mu \rangle \quad (8.13)$$

$$= \sum_{n=0}^{\infty} \langle c (dJ)_{\varphi^n} \cdot (d\varphi^n) f, \mu \rangle. \quad (8.14)$$

Since the center subspace is covariant,

$$\langle J, \partial_s \mu_s \rangle^c = \sum_{n=0}^{\infty} \langle c \frac{\beta_{\varphi^n}}{\beta} (dJ)_{\varphi^n} \cdot f_{\varphi^n}, \mu \rangle. \quad (8.15)$$

Using measure preservation, the center contribution is

$$\langle J, \partial_s \mu_s \rangle^c = \sum_{n=0}^{\infty} \langle \frac{c \circ \varphi^{-n}}{\beta \circ \varphi^{-n}} \beta dJ \cdot f, \mu \rangle. \quad (8.16)$$

It is clear from the above expression and Eq. 8.7 that the integrand is uniformly bounded. Secondly, we can evaluate the integral without explicitly propagating the tangent dynamics of the center subspace.

8.6 Updating the unstable contribution

We now turn our attention to the only remaining component of linear response yet to be discussed, the unstable contribution,

$$\langle J, \partial_s \mu_s \rangle^u := \sum_{k=0}^{\infty} \langle d(J \circ \varphi^k) \cdot (aq), \mu \rangle. \quad (8.17)$$

As with the center contribution, we have dropped the superscript on a , and this is justified for large run-up times of the regularized tangent equation (Eq. 8.6). With a sufficiently large run-up time, $a^0 \approx a$. To evaluate the unstable contribution (Eq. 8.17), we must perform integration by parts to regularize the integrand, which in Eq. 8.17, is exponentially growing in norm with k . The integration by parts is done similarly to in the case of uniformly hyperbolic maps (section 7.4.4). We first disintegrate the SRB measure on a partition of M subordinate to the unstable manifold (Ξ in Chapter 7). Then, we integrate by parts along each local unstable manifold. This yields the following regularized expression for the unstable contribution,

$$\langle J, \partial_s \mu_s \rangle^u := - \sum_{k=0}^{\infty} \langle (J \circ \varphi^k) (a g + b), \mu \rangle. \quad (8.18)$$

Here, as in Chapter 7, we have introduced a scalar field $b := \partial_\xi a$ to denote the unstable derivative of a . Recall that the distribution g is the unstable derivative of the logarithm of the SRB density conditioned on the unstable manifolds (Eq. 7.55).

8.6.1 The existence of unstable derivatives of a and c

The proof of existence of b is identical to that in the uniformly hyperbolic map case. Recall that in uniformly hyperbolic maps, we showed that the unstable derivative of v exists, and this existence implied the existence of b . The same holds in the case of uniformly hyperbolic flows that we consider here. From Eq. 8.8, we can see that, if v is differentiable on the unstable manifold, so is a since all other quantities on the right hand side of Eq. 8.8 are differentiable on the unstable manifold.

It remains to show that the regularized tangent vector field v is differentiable on the unstable manifold. To show that v in the uniformly hyperbolic map case has an unstable derivative, we defined a sequence $\{y^n\}$ as the unstable derivatives of the sequence $\{v^n\}$ whose limit is v . Here we follow the same argument so that y^n is defined as the sequence of unstable derivatives of v^n from Lemma 9. That is, writing $v^{n+} := d\varphi v^n + \chi$,

$$y^{n+1} := \partial_\xi v^{n+1} = \mathcal{P} \partial_\xi v^{n+} + (\partial_\xi \mathcal{P}) v^{n+}. \quad (8.19)$$

Before we show that the above sequence $\{y^n\}$ converges uniformly, we prove that it is well-defined for any bounded vector field sequence $\{v^n\}$ that is differentiable in the unstable direction. If $\{v^n\}$ has unstable derivatives at each n , this means that the first term in Eq. 8.19 is well-defined. Hence, we are left to prove that the operator \mathcal{P} is differentiable in the unstable direction so that the second term in Eq. 8.19 is well-defined because then, so is the first term, whose existence can be shown inductively from the differentiability of v^0 (by assumption) and differentiability of \mathcal{P} in the unstable direction.

Lemma 10. $\partial_\xi \mathcal{P}$ is a bounded linear operator.

Proof. First we show that \mathcal{P} is differentiable in the unstable direction. Recall that \mathcal{P} is an orthogonal projection out of the subspace $E^u \oplus E^c$. Thus, at every $x \in M$, it has the following matrix form,

$$\mathcal{P}_x = I - Q_x (Q_x^T Q_x)^{-1} Q_x^T, \quad (8.20)$$

where $Q_x = [q_x \ f_x]$ is an $m \times 2$ matrix and I is the $m \times m$ identity matrix. Moreover, $Q_x^T Q_x$ is a symmetric matrix with eigenvalues $1 - \theta_x$ and $1 + \theta_x$. Thus, using the matrix 2-norm, $\|(Q_x^T Q_x)\| = \max\{|1 - \theta_x|, |1 + \theta_x|\} > 0$, as long as $\theta_x = q_x \cdot f_x \neq 1$. This condition is indeed satisfied by uniformly hyperbolic flows wherein tangencies between the stable, unstable and center spaces are precluded.

Thus, $Q_x^T Q_x$ is invertible at all $x \in M$. Note that Q_x , and hence $Q_x^T Q_x$, are

differentiable in the unstable direction. Since matrix inverse is a differentiable function on the set of invertible matrices, $(Q_x^T Q_x)^{-1}$ is differentiable in the unstable direction. Since the right hand side of Eq. 8.20 is a composition of differentiable functions, \mathcal{P}_x is differentiable in the unstable direction. The linearity of $\partial_\xi \mathcal{P}$ follows from the definition of \mathcal{P} , Eq. 8.20.

Having shown that $\partial_\xi \mathcal{P}_x$ exists at all $x \in M$, we prove its uniform boundedness. By Chapter 7, Lemma 4, the unstable derivative of q , w is a bounded vector field. The vector field f has bounded derivatives in all directions by assumption. Hence, Q , Q^T and $Q^T Q$ have bounded unstable derivatives. Hence, to show that the derivative with respect of ξ of Eq. 8.20 is bounded, we only need to show that the unstable derivative of $(Q^T Q)^{-1}$ is bounded. But, $\partial_\xi (Q^T Q)^{-1} = (Q^T Q)^{-1} \partial_\xi (Q^T Q) (Q^T Q)^{-1}$, and $\partial_\xi (Q^T Q)$ is bounded. Further, $(Q^T Q)^{-1}$ is also bounded because $\|(Q^T Q)^{-1}\| = \max\{1/|1 - \theta|, 1/|1 + \theta|\} < \infty$. Thus, $\partial_\xi (Q^T Q)^{-1}$ is bounded. \square

Lemma 11. *The sequence y^n converges uniformly.*

Proof. Using the above lemma, we have shown that the sequence y^n , satisfies the following recursive equation (same as Eq. 8.19 with v^{n+} expanded),

$$y^{n+1} = \mathcal{P} \frac{d\varphi y^n}{\alpha} + \mathcal{P} \left(\frac{d^2\varphi(q, v^n)}{\alpha} + d\chi q \right) + \partial_\xi \mathcal{P}(d\varphi v^n + \chi). \quad (8.21)$$

To show that this sequence converges uniformly, we use Lemma 6 from Chapter 7.

To apply this lemma, we need to show that

$$\zeta^{n+1} := \mathcal{P} \left(\frac{d^2\varphi(q, v^n)}{\alpha} + d\chi q \right) + \partial_\xi \mathcal{P}(d\varphi v^n + \chi) \quad (8.22)$$

is uniformly exponentially converging. Subtracting ζ^n from ζ^{n+1} , and using the linearity of $d^2\varphi$, we get,

$$\zeta^{n+1} - \zeta^n = \frac{\mathcal{P}}{\alpha} d^2\varphi(q, v^n - v^{n-1}) + \partial_\xi \mathcal{P} d\varphi (v^n - v^{n-1}). \quad (8.23)$$

Since \mathcal{P}_x is an orthogonal projection operator with unit norm at each $x \in M$, and as we have shown in Lemma 10, $\|(\partial_\xi \mathcal{P})\|$ is a bounded operator, let $A > 0$ be such that

$\|(\partial_\xi \mathcal{P})_x\| < A$, for all $x \in M$. Then, for any $n \in \mathbb{Z}^+$,

$$\|\zeta^{n+1} - \zeta^n\| \leq \frac{\|d^2\varphi\|}{\alpha} \|v^n - v^{n-1}\| + A\|d\varphi\| \|v^n - v^{n-1}\| \quad (8.24)$$

$$\leq c\lambda^n, \quad (8.25)$$

for some $c > 0$, using uniform exponential convergence of $\{v^n\}$ in Lemma 9. Thus, ζ^n is uniformly exponentially converging, and hence Lemma 6 in Chapter 7 applies. Hence, y^n is a uniformly converging sequence, and the convergence, as per Lemma 6 of Chapter 7, is exponential. \square

Now since $\{v^n\}$ and its unstable derivative sequence $\{y^n\}$ have been shown to be uniformly converging, we have the relationship, $y = \partial_\xi v$, where y and v are the limits of the sequences $\{y^n\}$ and $\{v^n\}$. Finally, recall that we established this relationship, and showed that the regularized tangent vector field v is differentiable on the unstable manifold, en route to our main result: the differentiability in the unstable manifold of the scalar field a . Thus, we have now proven the existence of b , the unstable derivative of a .

Additionally, using the same argument, we have established that the scalar field c , is also differentiable on the unstable manifold. To wit, using Eq. 8.7, the differentiability of v in the unstable direction, combined with the fact that all other objects on the right hand side of Eq. 8.7 have unstable derivatives as well, it is clear that c is differentiable in the unstable direction.

8.7 Differentiating the unstable projection

The unstable contribution (Eq. 8.18) is modified from that in uniformly hyperbolic maps that we dealt with before because, since v is modified, a and consequently, b , are modified. By assumption, f is differentiable in all directions; let r denote the unstable derivative of f . Recall that the scalar field $b := \partial_\xi a$ is the unstable derivative of the scalar field a and its existence has been shown in the previous section (section 8.6.1). We also recall that $\alpha = \|d\varphi q\|$ is the expansion factor in the unstable direction and

that w is the unstable derivative of q . In section 8.6.1, we showed that the unstable derivative of v exists. Thus, it is valid to differentiate the expression for v (Eq. 8.6) in the unstable direction to yield,

$$y = \frac{d\varphi y + d^2\varphi(v, q)}{\alpha} + d\chi q - a w - b q - c r - \kappa f, \quad (8.26)$$

where, besides $r := \partial_\xi f$, we have introduced the scalar field $\kappa := \partial_\xi c$. Note that the existence of κ has also been proved in section 8.6.1. There are three unknowns, y , κ and b , in the above equation. Two more equations to close the above system can be obtained by differentiating the constraints, $v \cdot q = 0$ and $v \cdot f = 0$, which are true by construction (Lemma 9). Explicitly, differentiating these constraints in the unstable direction, we get

$$y \cdot q + v \cdot w = 0 \quad (8.27)$$

$$y \cdot f + v \cdot r = 0. \quad (8.28)$$

We set $y_x = 0$ for some μ -typical point x . Then, iterating Eq. 8.26 along the trajectory of x , and choosing b and κ so as to satisfy Eq. 8.27 and Eq. 8.28, we obtain y , b and κ along the trajectory of x . We can obtain expressions equivalent to the constraints, Eq. 8.27 and Eq. 8.28, if we directly differentiate the explicit expressions for a (Eq. 8.8) and c (Eq. 8.7) in the unstable direction. For completion, we note the unstable derivatives,

$$b = \frac{2\theta\eta}{(1-\theta^2)^2} v^+ \cdot (q - \theta f) + \frac{1}{1-\theta^2} \left(v^+ \cdot (w - \eta f - \theta r) + y^+ \cdot (q - \theta f) \right), \quad (8.29)$$

where

$$\begin{aligned} \eta &:= \partial_\xi \theta = r \cdot q + f \cdot w \\ y^+ &:= \partial_\xi v^+ = \frac{d^2\varphi(q, v) + d\varphi y}{\alpha} + d\chi q. \end{aligned} \quad (8.30)$$

Similarly, differentiating the explicit expression for c (Eq. 8.7) in the unstable direction,

$$\begin{aligned} \kappa &= \frac{2\theta\eta}{(1-\theta^2)^2} v^+ \cdot (f - \theta q) \\ &+ \frac{1}{1-\theta^2} \left(v^+ \cdot (r - \eta q - \theta w) + y^+ \cdot (f - \theta q) \right). \end{aligned} \quad (8.31)$$

8.8 The S3 algorithm for uniformly hyperbolic flows

We now collate the stable, unstable and center contributions, whose computation we discussed in the previous sections, to elucidate the overall S3 algorithm for uniformly hyperbolic flows. We reiterate that the algorithm below assumes that we have a one-dimensional unstable subspace spanned by $q_x \in T_x M$ and a one-dimensional center subspace, spanned by $f_x \in T_x M$, at every point $x \in M$.

Notation 5. *When a μ -typical phase point $x_0 \in M$ is fixed, the subscript n applied to a scalar function or a vector field, denoted h , is used to refer to its corresponding value at x_n . That is, when $x_0 \in M$ is fixed, $h_n := h_{\varphi^n x_0} = h_{x_n}$.*

Let K' be a sufficiently long run-up time such that the errors in v^0 , a^0 and c^0 , when compared to the corresponding limits of these sequences, v , a and c , are within a specified tolerance. The S3 algorithm is as follows:

1. Obtain a long primal trajectory $x_{-K'}, \dots, x_{N-1}$, where $x_{n+1} = \varphi x_n$, $-K' \leq n \leq (N-1)$, with $x_{-K'}$ chosen μ -a.e.
2. Obtain, at each point x_n , the unit tangent vector to the unstable manifold, q_n . The following procedure converges exponentially in n to the true value of q_n . Solve the homogeneous tangent equation with repeated normalization. That is, solve

$$\alpha_{n+1} q_{n+1} = (d\varphi)_n q_n, \quad n = -K', \dots, 0, 1, \dots, \quad (8.32)$$

with q_0 being set to a random vector in \mathbb{R}^m , and $\alpha_{n+1} = \|(d\varphi)_n q_n\|$.

3. Solve for v_n the following inhomogeneous tangent equation, which repeatedly projects v_n out of the center-unstable subspace,

$$v_{n+1} = (d\varphi)_n v_n + \chi_{n+1} - a_{n+1}q_{n+1} - c_{n+1}f_{n+1}, \quad n = -K', \dots, 0, 1, \dots, \quad (8.33)$$

where a_{n+1} and c_{n+1} are chosen such that $v_{n+1} \cdot q_{n+1} = 0$ and $v_{n+1} \cdot f_{n+1} = 0$. The initial condition $v_{-K'}$ is set to $0 \in \mathbb{R}^m$. This equation is also called the *regularized* tangent equation. The error in the solutions v_n when compared to the true value of the vector field v decreases exponentially with n (as shown in Lemma 9).

4. Solve the following second-order tangent equation for w_n , starting with $w_{-K'} = 0 \in \mathbb{R}^m$,

$$w_{n+1} = (1 - q_{n+1}q_{n+1}^T) \frac{(d^2\varphi)_n(q_n, q_n) + (d\varphi)_n w_n}{\alpha_{n+1}^2}, \quad n = -K', \dots, 0, 1, \dots, \quad (8.34)$$

The solutions w_n converge exponentially with n to the true values of the vector field w along the orbit (Chapter 7, Lemma 4).

5. Solve the following recursive second-order tangent equation for y_n , for each $n \in \mathbb{Z}^+$,

$$y_{n+1} = \frac{(d^2\varphi)_n(q_n, v_n) + (d\varphi)_n y_n}{\alpha_{n+1}} + (d\chi)_{n+1} q_{n+1} - b_{n+1}q_{n+1} - a_{n+1}w_{n+1} - c_{n+1}r_{n+1} - \kappa_{n+1}f_{n+1}, \quad n = -K', \dots, 0, \dots, \quad (8.35)$$

where the scalars b_{n+1} and κ_{n+1} is found using the relationships $y_{n+1} \cdot q_{n+1} = -v_{n+1} \cdot w_{n+1}$ and $y_{n+1} \cdot f_{n+1} = -v_{n+1} \cdot r_{n+1}$. The y_n , b_n and κ_n computed using this procedure also converge exponentially with n to their true values, as follows from Lemma 6.

6. Compute the unstable contribution as the following ergodic average:

$$\langle J, \partial_s \mu_s \rangle^u \approx -\frac{1}{N} \sum_{k=0}^{K-1} \sum_{n=0}^{N-1} J_{n+k} (a_n g_n + b_n). \quad (8.36)$$

7. Compute the center contribution as the following ergodic average:

$$\langle J, \partial_s \mu_s \rangle^c \approx \frac{1}{N} \sum_{k=0}^{K-1} \sum_{n=0}^{N-1} (dJ)_{n+k} \cdot f_{n+k} \frac{\beta_{n+k} c_n}{\beta_n}. \quad (8.37)$$

8. Compute the stable contribution as the following ergodic average:

$$\langle J, \partial_s \mu_s \rangle^s \approx \frac{1}{N} \sum_{n=0}^{N-1} (dJ)_n \cdot v_n. \quad (8.38)$$

9. The output of the S3 algorithm is the sum of the stable, unstable and center contributions, $\langle J, \partial_s \mu_s \rangle^u + \langle J, \partial_s \mu_s \rangle^c + \langle J, \partial_s \mu_s \rangle^s$. In the limit $N \rightarrow \infty$ and $K \rightarrow \infty$, in that order, the output converges to the true value of linear response (Ruelle's formula).

Remark 11. *Note that Ruelle's formula [235] does not assume that the flow of f is mixing. However, the integrability of the contribution to the sensitivity from the perturbation $\chi^u + \chi^c$, which is different from our center-unstable contribution (see below Theorem A of [235]), is assumed. We make an analogous time-discrete assumption on the summability of the center-unstable contribution*

$$\begin{aligned} & |\langle J, \partial_s \mu_s \rangle^c + \langle J, \partial_s \mu_s \rangle^u| \leq \sum_{k=0}^{\infty} |\langle d(J \circ \varphi^k) \cdot (ag + cf), \mu \rangle| \\ & = \sum_{k=0}^{\infty} \left| \langle (dJ \cdot f) \circ \varphi^k \frac{\beta \circ \varphi^k c}{\beta}, \mu \rangle - \langle J \circ \varphi^k (ag + b), \mu \rangle \right| \\ & < \infty. \end{aligned} \quad (8.39)$$

Remark 12. *The above summability assumption is needed for the existence of the center and unstable contributions. The sums may have subexponential decay, since*

for Hölder observables, which appear in the correlations above, correlations do not decay exponentially in general ([176] proves an exception in the case of smooth contact Anosov flows). In other words, unlike in uniformly hyperbolic diffeomorphisms considered in Chapter 7, the unstable and center contributions, are not, in general, evaluated accurately with the first few terms of the series.

8.9 Numerical example: Lorenz '63 system

The Lorenz'63 system is a three-variable reduced-order model of atmospheric convection [179]. It is a famous example of a chaotic system that first suggested that the earth's climate system is chaotic. The three variables x , y and z together capture thermal convection occurring with two-dimensional motions of a fluid kept between two parallel plates maintained at two different, constant temperatures [179]. The Lorenz'63 system of equations is given by

$$\frac{dx}{dt} = -s_1x + s_1y \tag{8.40}$$

$$\frac{dy}{dt} = x(s_2 - z) - y \tag{8.41}$$

$$\frac{dz}{dt} = xy - s_3z. \tag{8.42}$$

The standard parameter values are $s_1 = 10$, $s_2 = 28$ and $s_3 = 8/3$. Lorenz [179] and Saltzman [238] gave a physical interpretation of each coordinate function of the three-dimensional state vector, $x \equiv [x, y, z]^T$. All three are obtained from a truncated Fourier expansion of the stream function and the temperature deviation from a temperature profile that is linear between the two plates (this is an unstable fixed point of the dynamics corresponding to no fluid flow). The third coordinate z , in particular, is the only one with a nonzero ergodic average, and is proportional to the temperature deviation [179]. Note that the Lorenz'63 system has three distinct Lyapunov exponents: one positive at ≈ 0.9 , one zero, and one negative ≈ -14.6 . The center manifold, tangent to the flow direction, is one-dimensional. Our discrete time approximation to the flow, φ is a second-order Runge-Kutta numerical solver with




Figure 8-1: The parametric derivatives $d_{s_2}\langle z \rangle$ (left) of the Lorenz system with error bars indicating one standard deviation. On the right, $d_{s_2}\langle (z - 28)^2 \rangle$ computed using S3 at various different s_2 values on the x -axis are shown as black lines. The blue points on the right hand side subfigure are computed as ergodic averages along trajectories of length 1.6 billion (3.2 million time units). The other two parameters s_1 and s_3 are held at their standard values.

timestep size 0.002. Although the center direction of φ , the numerical approximation of the Lorenz system (Eq. 8.40), is not exactly equal to the flow direction, assuming they are equal leads to good numerical results for the overall sensitivity at this small timestep size.

The ensemble average $\langle z \rangle$ is known in the literature [168] to show an almost linear variation, with a slope of approximately 1, in the parameter s_2 . To validate the S3 algorithm, we compute this derivative $d_{s_2}\langle z \rangle$ and $d_{s_2}\langle (z - 28)^2 \rangle = 2d_{s_2}\langle z^2 + 784 - 56z \rangle = 2d_{s_2}\langle z^2 \rangle - 112\langle z \rangle$. As shown on the left hand side of Figure 8-1, the value of $d_{s_2}\langle z \rangle$ obtained by implementing the S3 algorithm (section 8.8) is close to the known value of 1 for a range of s_2 values. The S3 linear response, shown as blue crosses, with the blue lines indicating one standard deviation from the sample mean of 16 runs, is calculated along an orbit of length 3.2 million (6400 time units). The results of a second experiment with the objective function $(z - 28)^2$ are shown on the right hand side subfigure of Figure 8-1. We again use 3.2 million length orbits to calculate sensitivities using S3 (section 8.8), which are linearly extrapolated and shown as black lines. These lines closely trace the response curve $\langle (z - 28)^2 \rangle$ -vs.- s_2 computed by long-time averages (over 1.6 billion length orbits), shown as blue dots. For both objective functions, the center contribution is significant and thus, the correctness of the S3 sensitivities validates the S3 algorithm for uniformly hyperbolic flows.

8.10 Conclusion

In uniformly hyperbolic flows, the S3 algorithm we introduced in the previous chapter needs to be modified to account for the presence of the center direction. We modify the regularized tangent equation to be perpendicular to both the unstable direction as well as the center direction. This results in two scalar fields, a and c , which capture the history of projections of the tangent solutions on the unstable and center directions respectively. We derive the modified stable and unstable contributions, and the new center contribution to the overall sensitivity. The stable contribution is, as before, given as an ensemble average, $\langle dJ \cdot v \rangle$, where v is the regularized tangent solution that belongs to $(E^u \oplus E^c)^\perp$.

The unstable contribution is a sum of time-correlations between J and $ag + b$, where b is, as before, defined as the unstable derivative of a . While this expression is the same as in the uniformly hyperbolic map case, here the scalar fields a and consequently b are modified in the presence of the center direction. Furthermore, the decay of this time series is not exponential in general. We prove that the new regularized tangent vector field, v , is differentiable in the unstable direction and elucidate that this implies the existence of b , and that of κ , the unstable derivative of c . We then illustrate how to efficiently evaluate b from a second-order tangent equation that yields the unstable derivative of v .

As in uniformly hyperbolic maps, we see that the S3 algorithm is an ergodic-averaging scheme in which the individual components – the stable, unstable, and the new, center contributions – can be efficiently computed using information along a μ -typical trajectory.

Chapter 9

Remarks on the S3 algorithm

In this chapter, we comment on the error convergence of the S3 algorithm detailed in the last chapter, and its computational cost. First we describe a second numerical experiment (aside from the Baker's map considered in Chapter 7) to study the performance of the S3 algorithm (section 7.3.1).

9.1 Perturbations of the Solenoid Map

The Smale-Williams Solenoid map is a classical example of uniform hyperbolicity on a solid torus in \mathbb{R}^3 . We introduce a class of perturbed Solenoid maps expressed below in cylindrical coordinates:

$$\varphi_s([r, \theta, x^{(3)}]^T) = \begin{bmatrix} s_1 + (r - s_1)/4 + \cos \theta/2 \\ (2\theta + (s_2/4) \sin(4\theta)) \bmod (2\pi) \\ x^{(3)}/4 + \sin \theta/2 \end{bmatrix}. \quad (9.1)$$

The map can be expressed in Cartesian coordinates, denoted $[x^{(1)}, x^{(2)}, x^{(3)}]^T$, by left and right compositions of φ_s as in Eq. 9.1 with the following coordinate transformation and its inverse respectively: $x^{(1)} = r \cos \theta$, $x^{(2)} = r \sin \theta$. In Figure 9-1, we show the effect of parameter variation on the Solenoid attractor, which is shown on the $x^{(1)}$ - $x^{(2)}$ plane. The subfigures show the points on an orbit of length 2 million at the




Figure 9-1: The projection of the Solenoid attractor on the $x^{(1)}$ - $x^{(2)}$ plane at $s = [1, 0]^T$ (left), $s = [1, 1]^T$ (center) and $s = [2, 0]^T$ (right) respectively.

corresponding values of s . On the left is the attractor of the unperturbed map at the standard value of $s = [1, 0]^T$. The center subfigure, at the value $s = [1, 1]^T$, shows a change in the *shape* of the attractor due to an s_2 perturbation, while the rightmost subfigure indicates a change in the *size* of the attractor due to an s_1 perturbation. Qualitatively, we can also note change in the distribution of points on the perturbed attractor. For instance, the effect of an increase in s_2 is to reduce the probability of visiting the outer rim of the attractor. Both these types of qualitative changes, i.e., in the position and shape of, and in the distribution on the attractor are indicative of changes in the SRB measure and are captured precisely by the computation of Ruelle's formula.

Recall from Chapter 5 (or by inspecting Eq. 9.1) that the r and $x^{(3)}$ coordinate directions are stable at each point. The unstable direction, however, is not aligned with $\hat{\theta}$ everywhere. Thus, an s_1 perturbation corresponds to a stable vector field $\chi \in E^s$, and an s_2 perturbation has both stable and unstable components, albeit a small stable component.

In Chapter 5, we argued that the unstable contribution to the sensitivity is 0 for an s_1 perturbation because Ruelle's formula was split by decomposing χ as $\chi^u + \chi^s$. By contrast, in the S3 algorithm, although the two terms of the split Ruelle's formula are called stable and unstable contributions, the perturbation vector field χ is not split along E^s and E^u . Hence, in this case, both contributions are non-zero for both types of parameter perturbations although for the s_1 perturbation, the stable contribution




Figure 9-2: The sensitivities of $\langle J \rangle = \langle x^{(1)} \cdot x^{(1)} + x^{(2)} \cdot x^{(2)} \rangle$ with respect to the parameter s_1 (left) and s_2 (right) computed by S3 are shown as black lines at a number of parameter values on the $\langle J \rangle$ vs parameter curve.

dominates, and in the s_2 , the unstable dominates.

It is numerically verified that the S3 algorithm computes the correct sensitivity in both cases, as shown in Figure 9-2. The objective function, J is arbitrarily taken to be $J = r^2 = x^{(1)2} + x^{(2)2}$, and the sensitivity of $\langle J \rangle$ is computed at a range of values of s_1 (s_2) in the left (right) of Figure 9-2. The computed S3 sensitivities closely trace the curves $\langle J \rangle$ -vs- s (shown as blue dots) that are computed from ergodic averages using orbits of length 160 million. The S3 sensitivities, which are linearly extrapolated and shown as black lines, are computed along orbits of length 3.2 million at each value of s . The source code to generate S3 results on the Solenoid map can be found at [59].

9.2 Finite sample error analysis of the S3 algorithm

Recall the stable contribution from Chapter 7,

$$\langle J, \partial_s \mu_s \rangle^s = \langle dJ \cdot v, \mu \rangle, \quad (9.2)$$

which is practically computed as an ergodic average, $(1/N) \sum_{n=0}^{N-1} (dJ)_{x_n} \cdot v_{x_n}$, for an x_0 chosen μ -a.e. Since the convergence of the computed vector fields v^n to v is exponential, we can assume, for a sufficiently long run-up time, that the computed values of v along the orbit of x_0 are arbitrarily close to their true values. Thus, we only focus on quantifying the error, e^s , in the stable contribution due to a finite number of samples, N , used in the computation of the ergodic average.

In uniformly hyperbolic systems, the almost sure invariance principle, and its corollaries, the central limit theorem and the law of the iterated logarithm, apply to ergodic averages of Hölder continuous observables [84, 197]. Since both dJ , by assumption, and v , by Lemma 3, are Hölder continuous vector fields, $dJ \cdot v$ is also Hölder continuous. By applying the law of the iterated logarithm, we can obtain an almost sure upper bound for the error in the stable contribution. There exist a $\sigma^s > 0$ and an $N_0 \in \mathbb{N}$ such that, for almost every x , $N > N_0$ implies

$$e^s(x; N) = \left| \langle J, \partial_s \mu_s \rangle^s - (1/N) \sum_{n=0}^{N-1} (dJ)_{x_n} \cdot v_{x_n} \right| \leq \sigma^s \frac{\sqrt{2 \log \log N}}{\sqrt{N}}. \quad (9.3)$$

Now recall the unstable contribution from Chapter 7:

$$\langle J, \partial_s \mu_s \rangle^u = - \sum_{n=0}^{\infty} \langle J \circ \varphi^n(a g + b), \mu \rangle, \quad (9.4)$$

which, in practice, we compute for a finite number of terms K and a length of orbit N as

$$- \sum_{k=0}^{K-1} \frac{1}{N} \sum_{n=0}^{N-1} J_{k+n}(a_n g_n + b_n). \quad (9.5)$$

Analogous to the discussion above on the stable contribution, the functions g , a and b all exhibit exponential convergence in the length of the orbit used (see Lemmas 8, 5 and 7 respectively). Hence, we may assume that the values a_n, g_n and b_n computed by the algorithm are close to their corresponding true values for all $n \geq 0$, when a sufficiently long run-up before the estimation of Eq. 9.5 is used. Thus, we only explicitly discuss the error in the unstable contribution due to a finite N and K .

That is, we now characterize the error when computing the unstable contribution through Eq. 9.5 along an orbit of length $\mathcal{O}(N + K)$, which excludes the run-up time to reduce errors in the values of a, g and b to within machine precision. The error in

the unstable contribution,

$$e^u(x; N) \leq \left| \sum_{k \geq K} \langle J \circ \varphi^k(ag + b), \mu \rangle \right| + \sum_{k=0}^{K-1} \left| \langle J \circ \varphi^k(ag + b), \mu \rangle - \frac{1}{N} \sum_{n=0}^{N-1} J_{n+k}(a_n g_n + b_n) \right|. \quad (9.6)$$

Since $\langle (ag + b), \mu \rangle = 0$ (Theorem 3.1(b) of [232]), and from exponential decay of correlations of the Hölder continuous functions J and $ag + b$, the first term on the right hand side of (9.6), which is the bias in the S3 estimate (Eq. 9.5), is $\mathcal{O}(e^{-\gamma K})$, for some $\gamma > 0$. To bound the second term, we can again invoke the law of the iterated logarithm [84, 197]. There exist a $\sigma^u > 0$ and an $N_0 \in \mathbb{N}$ such that for almost every x , $N > N_0$ implies

$$e^u(x; N) = \left| \sum_{k=0}^{K-1} \frac{1}{N} \sum_{n=0}^{N-1} J_{k+n}(a_n g_n + b_n) + \langle J, \partial_s \mu_s \rangle^u \right| \leq \sigma^u K \frac{\sqrt{2 \log \log N}}{\sqrt{N}} + C e^{-\gamma K} \quad (9.7)$$

Invoking the exponential decay of correlations, the bias in the S3 estimate (sum of terms of Ruelle's formula for $k \geq K$ in Eq. 9.4) will typically be small even for small values of K , which can therefore be chosen independently of N to restrict the bias. Then, for a choice of constant K independent of N , the error in the unstable contribution is $\mathcal{O}(\sqrt{\log \log N}/\sqrt{N})$. Thus, from Eq. 9.3 and Eq. 9.7, the overall error in the S3 algorithm is $\mathcal{O}(\sqrt{\log \log N}/\sqrt{N})$. In Figure 9-3, a comparison of the error convergence of S3 against a naïve finite difference calculation is shown using the Solenoid map (section 9.1). To generate the results shown, s_1 is fixed at 1, and the derivative of $\langle J \rangle = x^{(1)}.x^{(1)} + x^{(2)}.x^{(2)}$ with respect to s_2 is computed at $s = [1, 0]^T$. The finite difference approximation calculated with N samples is as follows,

$$\langle J, \partial_s \mu_s \rangle_{\text{FD}} = \frac{\langle J \rangle_N([s_1, s_2 + \delta s]) - \langle J \rangle_N([s_1, s_2 - \delta s])}{2\delta s}, \quad (9.8)$$

where $\langle J \rangle_N(s)$ is an N -sample Monte Carlo estimate of $\langle J \rangle(s)$. Such an estimate can be computed either as an ergodic average along a single trajectory initialized




Figure 9-3: Comparison of relative error in the sensitivity of $\langle J \rangle$ to s_2 at $s = [1, 0]$ of the Solenoid map (section 9.1) as a function of the number of samples, between S3 and finite difference. The relative error is computed with respect to the true value, which is taken to be the S3 derivative obtained with $N = 9.6$ billion samples. The finite difference in the parameter s_2 used to calculate the derivatives is set to 0.001 in order to obtain the red dots. The green triangles are the derivatives computed using the S3 algorithm described in section 7.3.1. The dotted lines have a slope of -0.5.

Lebesgue almost everywhere, or as a sample average at samples according to μ on the attractor. To generate these plots efficiently, we do not compute ergodic averages along a single trajectory, as such a computation is a serial operation. Rather we sample average along multiple short trajectories in parallel, where each trajectory is initialized with a sample according to μ , which is achieved after a sufficiently long run-up time, starting from Lebesgue a.e. After this, each trajectory, with sufficiently long skips to allow correlation decay, can, in parallel, provide independent samples according to μ . This embarrassingly parallel operation is performed on a GPU (see [64] for the source code) and can compute Monte Carlo estimates using 10 billion samples in a few seconds. We observe in Figure 9-3 that the numerator of Eq. 9.8 is still overwhelmed by statistical noise, which follows the central limit theorem. As a result, the errors in the finite difference decay as $\mathcal{O}(1/\sqrt{N})$, which can be seen from the data points being well-approximated by a line with slope -0.5 (shown as a dotted red line). However, the central difference approximation grows as $\mathcal{O}(1/\delta s)$. As a result, the finite difference derivatives have much larger errors – in this case, two orders of magnitude larger errors – than the S3 derivatives (which are shown as green triangles), for the same number of samples. For instance, to achieve a 10% relative error, a central finite difference with $\delta s = 0.001$ requires more than a billion samples while the S3 algorithm shows less than 10% error even at 10,000 samples.

Note that the true value to compute relative error is taken to be the S3 derivative obtained from a ~ 10 billion sample computation, which is verified against the slope of the response curve $\langle J \rangle$ -vs- s_2 at $s_2 = 0$. Finally, we note that the error convergence in S3 is as predicted by the analysis above. The dotted green line, which has a slope equal to -0.5, well approximates the S3 data points. This concurs with the results in this section that $|e^u + e^s| \sim \mathcal{O}(1/\sqrt{N})$, ignoring the $\sqrt{\log \log N}$ term which only has a negligible effect on the slope.

9.3 Finite sample cost analysis of S3

In this section, we discuss the algorithmic efficiency of the S3 algorithm for differentiating SRB measures with respect to parameters in uniformly hyperbolic maps (7). For simplicity, we first consider differentiating the statistics of a scalar objective function with respect to one parameter perturbation. We determine the time complexity of the algorithm as a function of i) number of samples, N , ii) number of terms of Ruelle’s formula computed in the unstable contribution, K , iii) dimension of the unstable subspace, d_u ($d_u = 1$ in Chapters 7 and 8) and iv) system dimension, d . Since the memory/space complexity depends on the implementation of every piece of the algorithm, starting from the nonlinear solver, first- and second-order tangent solvers, we restrict ourselves to determining the time complexity. It is worth noting that many CFD codes are highly optimized for performance on HPC clusters and often provide tangent solvers through automatic differentiation. For implementing custom tangent solvers (and second-order tangent solvers needed in S3) that utilize these solutions, when such solvers are not already available, we must in addition use checkpointing schemes [270] in order to make the S3 algorithm feasible for very high dimensional problems. The estimate of time complexity here is also rough because we do not assume any sparsity or structure in the Jacobian matrices and the second-order derivative of φ , which typically exists and is exploited in tensor products to obtain better scaling with dimension. Here we discount the cost of the nonlinear solver as it is considered a part of the input to the S3 algorithm.

First we analyze the S3 algorithm for one-dimensional unstable manifolds, which is developed in this thesis. This algorithm is iterative with an $\mathcal{O}(N + K)$ length loop over a μ -typical orbit obtained from the nonlinear solver. In each iteration, a number of scalar objects and vector objects are time-evolved. Focusing on the more time-consuming vector operations, we start by noting that one step of an inhomogeneous tangent equation (such as the regularized tangent equation 7.25, for v), implemented naïvely, involves a Jacobian matrix-vector multiplication, followed by vector addition, which is $\mathcal{O}(d^2)$ operations. Another type of first-order tangent equation we solve (e.g.

for q) is the homogeneous tangent equation, which involves only the Jacobian-vector product, and is again $\mathcal{O}(d^2)$ operations. The number of first-order tangent equations that must be solved for every iteration of the outer loop is 2: one for the vector field q and the other for v . In the case of a multi-dimensional unstable manifold (see section 10.1), $d_u > 1$, a set of d_u equations for the forward Lyapunov vectors must be solved instead of just one for q . Further, we must apply a QR factorization of the $d \times d_u$ matrix of solutions at every timestep (at every iterate of the loop) in order to find an orthogonal basis for the unstable subspace at that sample point. The computational cost of the QR algorithm, without exploiting any available structure or sparsity, is $\mathcal{O}(dd_u^2)$. Thus, the overall number of operations at every iteration that involve first-order tangent solvers, assuming $d_u \ll d$, is $\mathcal{O}(d^2d_u)$, which reduces to $\mathcal{O}(d^2)$ operations for one-dimensional unstable manifolds.

A second-order tangent equation (e.g., Eq. 7.39 for w , Eq. 7.19 for y) solved naïvely for one timestep, involves computing a bilinear form (a tensor-vector product to form a matrix, followed by a matrix-vector product), possibly matrix-vector multiplications and a vector addition, which is overall $\mathcal{O}(d^3)$ operations. In the 1D unstable manifold algorithm, we need to solve one second-order tangent equation for w and one for y (see section 7.5.4). However, in the higher-dimensional unstable manifold case (section 10.1), the number of second-order derivatives of the forward Lyapunov vectors increases to $\binom{d_u}{2}$; we also need d_u -many vector fields that replace the role of y : derivatives of v in each unstable direction. Note that in this multi-dimensional unstable manifold case, g is also a vector of length d_u , as opposed to a scalar, that must be estimated using these second-order solutions. The number of operations involving second-order tangent equations that must be performed at every iteration of the $N + K$ loop is $\mathcal{O}(d^3d_u^2)$.

Thus, comparing the number of operations involving first-order and second-order equations explicitly discussed above, in order to estimate the asymptotic time complexity, we only need to examine the number of second-order tangent evolutions in the S3 algorithm. In other words, the second-order tangent equations are the most expensive computation in the S3 algorithm.

The assembly of the stable contribution is straightforward since it is simply an ergodic average (Eq. 9.2) that adds d operations (to compute $dJ \cdot v$) per iteration. In order to assemble each of the K terms in the unstable contribution (Eq. 9.5), an N -sample ergodic average must be computed using all the vectors and scalars computed within the loop, which also adds $\mathcal{O}(d)$ operations per sample. However, K N -length orbits are available within an $N + K$ length orbit, obviating the need for an $\mathcal{O}(NK)$ loop. Thus, to obtain the stable and unstable contributions, we only need $\mathcal{O}((N + K)d^3d_u^2)$ operations. As discussed in section 9.2, since K is of constant order compared to N , the time complexity can be stated as $\mathcal{O}(Nd^3d_u^2)$.

9.3.1 Discrete adjoint of the S3 algorithm

Above we discussed the time complexity of the S3 algorithm to compute the derivative with respect to one system parameter of the expectation of one observable with respect to the SRB measure. When the number of observables is increased, the algorithm described in section 7.3.1 is efficient. The observable J enters only the final computation of the ergodic averages and hence the complexity of each iteration is increased (additively) by a factor of d that depends only on the number of observables. Hence, the time complexity of the algorithm does not change, and remains $\mathcal{O}(Nd^3d_u^2)$.

Now consider increasing the number of parameters to $d_p > 1$. We must compute one vector field v per parameter perturbation. Thus, the number of operations involving first-order tangent equations increases to $\mathcal{O}(d^2d_u + d^2d_p)$. The number of operations per iteration that involve second-order tangent equations also increases to $\mathcal{O}(d^3d_u^2 + d^3d_u d_p)$ since each first-order regularized tangent vector field v (and there are d_p such vector fields corresponding to d_p parameter perturbations) needs to be differentiated in d_u unstable directions. Thus, the time complexity of the S3 algorithm becomes $\mathcal{O}(Nd^3(d_u^2 + d_u d_p))$. When the number of parameters is much larger than d_u , this cost can be prohibitive and thus, it is more efficient to implement a discrete adjoint of the S3 algorithm.

That is, we isolate the steps of the S3 algorithm that depend on the parameters and apply reverse-mode differentiation to these steps. While the derivation of the

	Forward-mode (section 7.3.1)	Reverse-mode
# of 2^{nd} order tangents/adjoints	$d_u(d_u + 1)/2 + d_p d_u$	$d_u(d_u + 1)/2 + d_J$
Time complexity	$\mathcal{O}(Nd^3(d_u^2 + d_p d_u))$	$\mathcal{O}(Nd^3(d_u^2 + d_J d_u))$

Table 9.1: Summary of the cost of the S3 algorithm as described in section 7.3.1 and its reverse-mode/adjoint version.

discrete adjoint is beyond the scope of this thesis, we mention that the adjoint, y^* , of the d_p second-order tangent equations for y (Eq. 7.19) are now d_J first-order adjoint equations (involving the Jacobian transpose). Here d_J is the number of observables J . The adjoint of the regularized tangent equation (Eq. 7.25) is an adjoint equation that involves a second order derivative of φ :

$$v_n^* = (d^2\varphi_n q_n)^T y_n^* + (d\varphi)_n^T v_{n+1}^* + (dJ)_n. \quad (9.9)$$

Thus, the number of first-order adjoint equations in the adjoint S3 algorithm is $d_J \times d_u$, and the second-order adjoint equations, such as the one in Eq. 9.9, are d_J in number. Hence, the cost per timestep for the first-order equations is $\mathcal{O}(d^2 d_J d_u)$, and as usual, the second-order equations are the most expensive computation. Although the number of second-order equations does not depend on d_u , the cost of evolution is not independent of d_u . From Eq. 9.9, we can see that, in the case of higher-dimensional unstable manifolds, the number of $d^2\varphi$ -vector products is d_u . Hence, the cost per timestep of the second-order adjoint equations is $\mathcal{O}(d^3 d_J d_u)$. The cost of the second-order tangent equations, corresponding to the cross-derivatives and self-derivatives of the forward Lyapunov vectors still remains as in the tangent S3 algorithm, at $\mathcal{O}(d^3 d_u^2)$. Thus, the complexity of the adjoint algorithm that gives the derivatives of d_J expectations with respect to d_p parameters is $\mathcal{O}(Nd^3(d_u^2 + d_J d_u))$. When $d_J < d_u d_p$, the adjoint algorithm is more efficient than the forward-mode version in section 7.3.1, which, as described above, has a complexity of $\mathcal{O}(Nd^3(d_u^2 + d_u d_p))$. The number of second-order tangent equations, the most expensive computation in every iteration, and the time complexity of the S3 algorithm are summarized in Table 9.1.

Chapter 10

Summary and Future Work

The main contribution of this thesis is a new algorithm, called space-split sensitivity (S3), to compute linear response, or the sensitivities of statistics to parameters, in chaotic dynamical systems. The thesis presents the derivation of S3 with Ruelle's linear response formula [232][234] as the starting point and works exclusively with the simplest class of chaotic systems: uniformly hyperbolic systems with one-dimensional unstable manifolds. We first show that Ruelle's formula leads to an ill-conditioned direct computation (Chapter 4). To derive S3, we decompose this formula into a stable and unstable contribution, by splitting the parameter perturbation by means of solving a regularized tangent equation (Chapter 7). This leads to stable and unstable contributions to the sensitivity for both of which efficient methods of evaluation are developed. The resulting algorithm – the S3 algorithm – evaluates the stable and unstable contribution as ergodic averages that converge like a typical Monte Carlo integration, which has a dimension-independent error convergence rate of $\mathcal{O}(1/\sqrt{N})$ (ignoring a factor logarithmic in $\log N$) with N samples.

The ergodic-averaging computation of the unstable contribution involves novel recursive procedures to estimate i) the *density gradient* of the SRB density conditioned on the unstable manifold, and ii) the derivative of the one-dimensional unstable subspace in its own direction. The latter recursion is developed in Chapter 6. The density gradient is a fundamental object whose regularity we link, in a different study [248], to the validity of linear response [277][276].

In this thesis, we also discuss an application of an existing algorithm for sensitivity computation in chaotic systems, based on shadowing [206][211][273], to the problem of suppressing thermoacoustic instabilities in a reduced-order model of a gas turbine. We show, through examples (Chapter 3), that generally speaking, shadowing-based algorithms are not guaranteed to converge to the true sensitivity, which is consistent with the analysis done by Ni [209]. Owing to the low dimension of the unstable subspace compared to the overall dimension of the system, the shadowing sensitivities are accurate enough to be employed in optimization and data assimilation applications in the considered thermoacoustic example. As shown by Ni [209], shadowing only computes one component of Ruelle’s response – it disregards the change in the SRB probability density on the unstable manifold due to the parameter perturbation.

The S3 algorithm, on the other hand, is a provably convergent computation of linear response and has an efficient, scalable implementation. In Chapter 8, we extend the S3 algorithm from the discrete-time case presented in Chapter 7 to the continuous-time setting, which is representative of chaotic systems stemming from numerical solutions of ODEs. Before we close, we give a brief blueprint for the extension of S3 to problems with arbitrary-dimensional unstable manifolds.

10.1 S3 extension to multi-dimensional unstable manifolds

For the notations and definitions we shall use here, we refer the reader to Chapters 7 and 8. We briefly summarize the S3 decomposition and the algorithm (Chapter 7) before we discuss their multi-dimensional extensions. The S3 decomposition is achieved by orthogonalizing the conventional tangent solution with respect to the unstable subspaces along a long trajectory. This procedure yields a regularized tangent vector field v using which the scalar component, a , of the split parameter perturbation in the unstable direction, q , is determined.

The sensitivity to the unstable vector field so achieved – aq – is called the unstable

contribution. The stable contribution is the remaining linear response, and can be calculated just as tangent sensitivities are in non-chaotic systems, but using the regularized tangent solutions in place of the conventional tangent solutions. The unstable contribution is integrated by parts on the unstable manifold to yield an expression in which, unlike the original form of Ruelle’s formula, the integrand is uniformly bounded. Then, exploiting the exponential decay of correlations, this regularized formula can be computed fast as a sum of Monte Carlo integrals. However, the problem is, we obtain two unknown functions – the derivative of a on the unstable manifold, and the density gradient.

The density gradient is the unstable derivative of the logarithmic conditional density of the SRB measure on the unstable manifold. We take the unstable derivative of Pesin’s formula [148][169] to obtain a recursive evaluation of the density gradient along a typical trajectory. We show that starting this recursion with an arbitrary scalar field leads to an exponential convergence to the true density gradient. A recursive procedure is also derived for the unstable derivative of a by combining the unstable derivatives of v and q . The evaluation of the latter two unstable derivatives requires solving two second-order tangent equations, which are the most computationally intensive steps of the S3 algorithm (see section 9.3).

Overall, the S3 algorithm computes Ruelle’s formula as a well-conditioned ergodic average. An important direction of future work is to make S3 applicable to systems with higher-dimensional unstable manifolds. Instead of being one-dimensional, as we have assumed throughout Chapters 7 and 8, suppose E_x^u is an $m_u \geq 1$ dimensional subspace at each $x \in M$. We anticipate that the S3 algorithm then involves the following changes, which will be pursued in a future work:

1. We must compute, e.g. using Ginelli’s algorithm [114], an orthonormal basis for the unstable subspaces along an orbit. This more general procedure reduces to Step 2 of the S3 algorithm (section 5.6), in the case of a 1D unstable manifold. Let $q_{x_n}^i, 1 \leq i \leq m_u, 1 \leq n \leq N$ be the resultant orthogonal tangent vectors that span $E_{x_n}^u, 1 \leq n \leq N$.

2. The regularized tangent equation must now be orthogonalized with respect to the computed unstable subspaces along the trajectory, as opposed to with respect to a single direction. This orthogonalization results in a regularized tangent vector field v and m_u different scalar fields, $a^i, 1 \leq i \leq m_u$, in place of a single scalar field a . The stable contribution may be computed using v just as in the 1D case (Eq. 7.28).

3. The first step in treating the unstable contribution is the integration by parts on the unstable manifold. In the multi-dimensional unstable manifold case, the parameterization Φ^x is now a map from $[0, 1]^{m_u}$ to Ξ_x , such that $\nabla_\xi \Phi^{x_n}(0)$ maps the standard Euclidean basis vector e^i in \mathbb{R}^{m_u} to $q_{x_n}^i$. Using disintegration followed by integration by parts on the m_u -dimensional unstable manifold results in a regularized expression for the unstable contribution, analogous to the 1D case (Eq. 7.32),

$$\langle J, \partial_s \mu_s \rangle^u = - \sum_{n=0}^{\infty} \langle J \circ \varphi^n (a \cdot g + \text{div}^u a), \mu \rangle. \quad (10.1)$$

Here $a : M \rightarrow \mathbb{R}^{m_u}$ and $g : M \rightarrow \mathbb{R}^{m_u}$ are now vector-valued functions on M . The term $\text{div}^u a$ refers to the unstable divergence of the unstable perturbation field $\sum_i a^i q^i$. Now, analogous to the computation of $ag + b$ in the 1D unstable manifold, which we tackled in this thesis, we must derive a new computation for the derivative of $a \cdot g + \text{div}^u a$. Such a computation must similarly be recursive so that it can be efficiently carried out using information available along an orbit.

Since we used the unstable derivatives of α , v and q in the computation of $ag + b$ in the 1D unstable manifold case, we expect analogous derivatives are necessary in higher dimensions. In particular, the unstable derivatives – which are m_u -dimensional gradients, i.e., partial derivatives taken with respect to each ξ^i corresponding to q^i – of v and each q^i are needed. The higher dimensional analog of α are the diagonal elements of the R matrix in the QR-based iterative

computation in Ginelli's algorithm [114]; the unstable gradient of the diagonal elements of R are needed. Computing these required derivatives efficiently is the crux of this future direction that will enable the application of S3 to systems with arbitrary dimensional unstable manifolds.

From the outline sketched above, it is amply clear that the S3 algorithm presented here serves as an excellent starting point from which to extend to higher-dimensional unstable manifolds. Moreover, we have proved the existence of the S3 decomposition and the convergence of the algorithm in uniformly hyperbolic systems. The proof of convergence presented in this thesis is also readily extensible to the case of higher-dimensional unstable manifolds. Here we have shown that both the stable and unstable contributions to the sensitivity, computed as per the S3 algorithm, have an error convergence that declines as a Monte Carlo computation of an ergodic average, i.e., as $\mathcal{O}(1/\sqrt{N})$ along an orbit of length N . Thus, this work is the first step, with a clear path forward for efficient extensions to arbitrary dimensional systems, toward circumventing the poor convergence rate of Ruelle's formula for linear response.

Bibliography

- [1] Rafail V Abramov and Andrew J Majda. Blended response algorithms for linear fluctuation-dissipation for complex nonlinear dynamical systems. *Nonlinearity*, 20(12):2793–2821, oct 2007. doi: 10.1088/0951-7715/20/12/004. URL <https://doi.org/10.1088/0951-7715/20/12/004>.
- [2] R.V. Abramov and A.J. Majda. New approximations and tests of linear fluctuation-response for chaotic nonlinear forced-dissipative dynamical systems. *J Nonlinear Sci*, 18:303–341, 2008. doi: 10.1007/s00332-007-9011-9.
- [3] Alistair Adcroft, Jean-Michel Campin, Stephanie Dutkiewicz, Constantinos Evangelinos, David Ferreira, Gael Forget, Baylor Fox-Kemper and Patrick Heimbach, Chris Hill, Ed Hill, Helen Hill, Oliver Jahn, Martin Losch, John Marshall, Guillaume Maze, Dimitris Menemenlis, and Andrea Molod. Mitgcm user manual. http://mitgcm.org/public/r2_manual/latest/, 2008.
- [4] Harish Agarwal, John E Renaud, Evan L Preston, and Dhanesh Padmanabhan. Uncertainty quantification using evidence theory in multidisciplinary design optimization. *Reliability Engineering & System Safety*, 85(1):281–294, 2004.
- [5] Roger Alexander. Diagonally implicit runge–kutta methods for stiff ode’s. *SIAM Journal on Numerical Analysis*, 14:1006–1021, 1977. doi: 10.1137/0714068.
- [6] Mihai Alexe, Oleg Roderick, Mihai Anitescu, Jean Utke, Thomas Fanning, and Paul Hovland. Using automatic differentiation in sensitivity analysis of nuclear simulation models. *Transactions of the American Nuclear Society*, 102:235–237, 2010.
- [7] Grégoire Allaire. A review of adjoint methods for sensitivity analysis, uncertainty quantification and optimization in numerical codes. *Ingénieurs de l’Automobile*, 836:33–36, July 2015. URL <https://hal.archives-ouvertes.fr/hal-01242950>.
- [8] Brian Ancell and Gregory J. Hakim. Comparing adjoint- and ensemble-sensitivity analysis with applications to observation targeting. *Monthly Weather Review*, 135(12):4117 – 4134, 01 Dec. 2007. doi: 10.1175/2007MWR1904.1. URL <https://journals.ametsoc.org/view/journals/mwre/135/12/2007mwr1904.1.xml>.

- [9] Timothy Andrews, Jonathan M. Gregory, Mark J. Webb, and Karl E. Taylor. Forcing, feedbacks and climate sensitivity in cmip5 coupled atmosphere-ocean climate models. *Geophysical Research Letters*, 39(9), 2012. doi: <https://doi.org/10.1029/2012GL051607>. URL <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2012GL051607>.
- [10] D. V. ANOSOV. Geodesic flows and closed riemannian manifolds with negative curvature. *Proc. Steklov Inst. Math.*, 90, 1967.
- [11] Somchai Ansuaj, Farrokh Shokooh, and Roland Schinzinger. Parameter estimation for induction machines based on sensitivity analysis. *IEEE Transactions on Industry Applications*, 25(6):1035–1040, 1989.
- [12] Zin Arai. On hyperbolic plateaus of the h enon map. *Experimental Mathematics*, 16(2):181–188, 2007. doi: 10.1080/10586458.2007.10128992. URL <https://www.tandfonline.com/doi/abs/10.1080/10586458.2007.10128992>.
- [13] V. Ara ujo, I. Melbourne, and P. Varandas. Rapid Mixing for the Lorenz Attractor and Statistical Limit Laws for Their Time-1 Maps. *Commun. Math. Phys.*, 340(3):901–938, December 2015. ISSN 0010-3616, 1432-0916. doi: 10.1007/s00220-015-2471-0. URL <http://link.springer.com/10.1007/s00220-015-2471-0>.
- [14] V. Ara ujo, I. Melbourne, and P. Varandas. Rapid Mixing for the Lorenz Attractor and Statistical Limit Laws for Their Time-1 Maps. *Commun. Math. Phys.*, 340(3):901–938, December 2015. ISSN 0010-3616, 1432-0916. doi: 10.1007/s00220-015-2471-0. URL <http://link.springer.com/10.1007/s00220-015-2471-0>.
- [15] L. Arnold, H. Crauel, and J.-P. Eckmann. *Lyapunov Exponents*. Lectures Notes in Mathematics 1486. Springer-Verlag, Berlin Heidelberg, 1990. ISBN 3540665463. doi: 10.1007/978-3-540-78319-0.
- [16] Ludwig Arnold. *The Multiplicative Ergodic Theorem on Bundles and Manifolds*, pages 163–199. Springer Berlin Heidelberg, Berlin, Heidelberg, 1998. ISBN 978-3-662-12878-7. doi: 10.1007/978-3-662-12878-7_4. URL https://doi.org/10.1007/978-3-662-12878-7_4.
- [17] Tony Arts and Muriel Lambert De Rouvrot. Aero-thermal performance of a two dimensional highly loaded transonic turbine nozzle guide vane: A test case for inviscid and viscous flow computations. In *ASME 1990 International Gas Turbine and Aeroengine Congress and Exposition*, pages V001T01A106–V001T01A106. American Society of Mechanical Engineers, 1990. doi: 10.1115/1.2927978.
- [18] P. Ashwin and A.S. von der Heydt. Extreme sensitivity and climate tipping points. *J Stat Phys*, 179:1531–1552, 2020. doi: 10.1007/s10955-019-02425-x.

- [19] autodiff.org. Tools for automatic differentiation, 2020. URL <http://www.autodiff.org/?module=Tools>.
- [20] WJ Baars, N Hutchins, and I Marusic. Reynolds number trend of hierarchies and scale interactions in turbulent boundary layers. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 375(2089):20160077, 2017. doi: 10.1098/rsta.2016.0077.
- [21] Wael Bahsoun and Benoît Saussol. Linear response in the intermittent family: Differentiation in a weighted c^0 -norm. *Discrete & Continuous Dynamical Systems*, 36(12):6657–6668, 2016. doi: 10.3934/dcds.2016089.
- [22] Wael Bahsoun, Stefano Galatolo, Isaia Nisoli, and Xiaolong Niu. A rigorous computational approach to linear response. *Nonlinearity*, 31:1073–1109, 2018. doi: 10.1088/1361-6544/aa9a88.
- [23] Christopher TH Baker, GA Bocharov, CAH Paul, and FA Rihan. Computational modelling with functional differential equations: Identification, selection, and sensitivity. *Applied Numerical Mathematics*, 53(2-4):107–129, 2005.
- [24] Gabriel Balaban, Henrik Finsberg, Hans Henrik Odland, Marie E. Rognes, Stian Ross, Joakim Sundnes, and Samuel Wall. High-resolution data assimilation of cardiac mechanics applied to a dyssynchronous ventricle. *International Journal for Numerical Methods in Biomedical Engineering*, 33(11):e2863, 2017. doi: <https://doi.org/10.1002/cnm.2863>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/cnm.2863>. e2863 cmn.2863.
- [25] V. Baladi and M. Todd. Linear response for intermittent maps. *Commun. Math. Phys.*, 347:857–874, 2016. doi: 10.1007/s00220-016-2577-z.
- [26] V Baladi, J-P Eckmann, and D Ruelle. Resonances for intermittent systems. *Nonlinearity*, 2:119, 1989. doi: 10.1088/0951-7715/2/1/007.
- [27] Viviane Baladi. *Positive transfer operators and decay of correlations*, volume 16. World scientific, 2000. doi: 10.1142/3657.
- [28] Viviane Baladi. *Dynamical zeta functions and dynamical determinants for hyperbolic maps*. Springer, 2018. doi: 10.1007/978-3-319-77661-3.
- [29] Viviane Baladi and Daniel Smania. Linear response for smooth deformations of generic nonuniformly hyperbolic unimodal maps. *Annales scientifiques de l'ENS*, 45(6):861–926, 2012.
- [30] Pierre Baldi and Amir F Atiya. How delays affect neural dynamics and learning. *IEEE Transactions on Neural Networks*, 5(4):612–621, 1994.
- [31] J-FM Barthelemy and Laura E Hall. Automatic differentiation as a tool in engineering design. *Structural optimization*, 9(2):76–82, 1995. doi: 10.1007/BF01758823.

- [32] Ryne Beeson and N. Sri Namachchivaya. Particle filtering for chaotic dynamical systems using future right-singular vectors. *Nonlinear Dyn*, June 2020. ISSN 0924-090X, 1573-269X. doi: 10.1007/s11071-020-05727-y. URL <http://link.springer.com/10.1007/s11071-020-05727-y>.
- [33] C. Bischof, G. Corliss, L. Green, A. Griewank, K. Haigler, and P. Newman. Automatic differentiation of advanced cfd codes for multidisciplinary design. *Computing Systems in Engineering*, 3(6):625–637, 1992. ISSN 0956-0521. doi: [https://doi.org/10.1016/0956-0521\(92\)90014-A](https://doi.org/10.1016/0956-0521(92)90014-A). URL <https://www.sciencedirect.com/science/article/pii/095605219290014A>.
- [34] Christian H. Bischof, H. Martin Bücker, Paul Hovland, Uwe Naumann, and Jean Utke. *Advances in Automatic Differentiation*. Springer, Berlin, Heidelberg, 2008. doi: 10.1007/978-3-540-68942-3.
- [35] Christian H Bischof, Paul D Hovland, and Boyana Norris. On the implementation of automatic differentiation tools. *Higher-Order and Symbolic Computation*, 21(3):311–331, 2008. doi: 10.1007/s10990-008-9034-4.
- [36] Ron F Blackwelder and Leslie SG Kovaszny. Time scales and correlations in a turbulent boundary layer. *The Physics of Fluids*, 15(9):1545–1554, 1972. doi: 10.1063/1.1694128.
- [37] Patrick J Blonigan. Adjoint sensitivity analysis of chaotic dynamical systems with non-intrusive least squares shadowing. *Journal of Computational Physics*, 348:803–826, 2017. doi: 10.1016/j.jcp.2017.08.002.
- [38] Patrick J. Blonigan and Qiqi Wang. Probability density adjoint for sensitivity analysis of the mean of chaos. *Journal of Computational Physics*, 270:660 – 686, 2014. ISSN 0021-9991. doi: <https://doi.org/10.1016/j.jcp.2014.04.027>. URL <http://www.sciencedirect.com/science/article/pii/S0021999114002903>.
- [39] Patrick J. Blonigan, Pablo Fernandez, Scott M. Murman, Qiqi Wang, Georgios Rigas, and Luca Magri. Toward a chaotic adjoint for LES. *arXiv e-prints*, art. arXiv:1702.06809, February 2017.
- [40] Patrick J Blonigan, Qiqi Wang, Eric J Nielsen, and Boris Diskin. Least-squares shadowing sensitivity analysis of chaotic flow around a two-dimensional airfoil. *AIAA Journal*, pages 658–672, 2017. doi: 10.2514/1.J055389.
- [41] Patrick J. Blonigan, Mohammad Farazmand, and Themistoklis P. Sapsis. Are extreme dissipation events predictable in turbulent fluid flows? *Phys. Rev. Fluids*, 4:044606, Apr 2019. doi: 10.1103/PhysRevFluids.4.044606. URL <https://link.aps.org/doi/10.1103/PhysRevFluids.4.044606>.
- [42] R. S. Blumenthal, A. K. Tangirala, R. Sujith, and W. Polifke. A systems perspective on non-normality in low-order thermoacoustic models: Full

- norms, semi-norms and transient growth. *International Journal of Spray and Combustion Dynamics*, 9(1):19–43, 2016. ISSN 1756-8277. doi: 10.1177/1756827716652474. URL <http://scd.sagepub.com/lookup/doi/10.1177/1756827716652474>.
- [43] Daniel J Bodony and Sanjiva K Lele. Current status of jet noise predictions using large-eddy simulation. *AIAA journal*, 46(2):364–380, 2008. doi: 10.2514/1.24475.
- [44] T. Bomfim, A. Castro, and P. Varandas. Differentiability of thermodynamical quantities in non-uniformly expanding dynamics. *Advances in Mathematics*, 292:478–528, 2016. ISSN 0001-8708. doi: <https://doi.org/10.1016/j.aim.2016.01.017>. URL <https://www.sciencedirect.com/science/article/pii/S000187081600044X>.
- [45] Emanuele Borgonovo and Elmar Plischke. Sensitivity analysis: a review of recent advances. *European Journal of Operational Research*, 248(3):869–887, 2016.
- [46] Rufus Bowen, Jean-René Chazottes, and David Ruelle. *Equilibrium States and the Ergodic Theory of Anosov Diffeomorphisms*, volume 471. Springer Lecture Notes in Mathematics, 2008. doi: 10.1007/978-3-540-77695-6.
- [47] Vincent Brunet and Sébastien Deck. Zonal-detached eddy simulation of transonic buffet on a civil aircraft type configuration. *NOTES ON NUMERICAL FLUID MECHANICS AND MULTIDISCIPLINARY DESIGN*, 97:182, 2008.
- [48] H Martin Bücker, George Corliss, Paul Hovland, Uwe Naumann, and Boyana Norris. *Automatic differentiation: applications, theory, and implementations*, volume 50. Springer Science & Business Media, 2006.
- [49] Oliver Butterley and Carlangelo Liverani. Smooth anosov flows: correlation spectra and stability. *J. Mod. Dyn*, 1(2):301–322, 2007. doi: 10.3934/jmd.2007.1.301.
- [50] Tamás Bóday, Valerio Lucarini, and Frank Lunkeit. Can we use linear response theory to assess geoengineering strategies? *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 30(2):023124, 2020. doi: 10.1063/1.5122255. URL <https://doi.org/10.1063/1.5122255>.
- [51] Yong-Yan Cao and Paul M Frank. Analysis and synthesis of nonlinear time-delay systems via fuzzy control approach. *IEEE Transactions on fuzzy systems*, 8(2):200–211, 2000.
- [52] Jesse Capecelatro, Daniel J Bodony, and Jonathan B Freund. Adjoint-based sensitivity and ignition threshold mapping in a turbulent mixing layer. *Combustion Theory and Modelling*, pages 1–33, 2018. doi: 10.2514/6.2017-0846.

- [53] W. Castaings, D. Dartus, M. Honnorat, F.-X. Le Dimet, Y. Loukili, and J. Monnier. Automatic differentiation: A tool for variational data assimilation and adjoint sensitivity analysis for flood modeling. In Martin Bückner, George Corliss, Uwe Naumann, Paul Hovland, and Boyana Norris, editors, *Automatic Differentiation: Applications, Theory, and Implementations*, pages 249–262, Berlin, Heidelberg, 2006. Springer Berlin Heidelberg. ISBN 978-3-540-28438-3. doi: 10.1007/3-540-28438-9_22.
- [54] Massimo Cencini and Francesco Ginelli. Lyapunov analysis: from dynamical systems theory to applications. *Journal of Physics A: Mathematical and Theoretical*, 46(25):250301, jun 2013. doi: 10.1088/1751-8113/46/25/250301. URL <https://doi.org/10.1088/1751-8113/46/25/250301>.
- [55] B. Cessac. Linear response in neuronal networks: From neurons dynamics to collective response. *Chaos*, 29:103105, 2019. doi: 10.1063/1.5111803.
- [56] N. CHANDRAMOORTHY and Q. WANG. Supplementary materials for “on the probability of finding a nonphysical solution”. *Github*, v1.0, 2020. URL <https://github.com/nishaChandramoorthy/shadowing>.
- [57] Nisha Chandramoorthy. nishachandramoorthy/s3: Differential clv method. 1.0.0, July 2020. doi: 10.5281/zenodo.3941678. URL <https://doi.org/10.5281/zenodo.3941678>.
- [58] Nisha Chandramoorthy. nishaChandramoorthy/energies: Thermoacoustics paper, August 2020. URL <https://doi.org/10.5281/zenodo.3996268>.
- [59] Nisha Chandramoorthy. Space-split sensitivity for efficient linear response of chaos. <https://github.com/nishaChandramoorthy/linearResponse>, 2021.
- [60] Nisha Chandramoorthy and Qiqi Wang. Sensitivity computation of statistically stationary quantities in turbulent flows. In *AIAA Aviation 2019 Forum*, 2019. doi: 10.2514/6.2019-3426. URL <https://arc.aiaa.org/doi/abs/10.2514/6.2019-3426>.
- [61] Nisha Chandramoorthy and Qiqi Wang. On the probability of finding a non-physical solution through shadowing. *arXiv e-prints*, art. arXiv:2010.13768, October 2020.
- [62] Nisha Chandramoorthy and Qiqi Wang. A computable realization of Ruelle’s formula for linear response of statistics in chaotic systems. *arXiv e-prints*, art. arXiv:2002.04117, February 2020.
- [63] Nisha Chandramoorthy and Qiqi Wang. An ergodic averaging method to differentiate covariant Lyapunov vectors. *arXiv e-prints*, art. arXiv:2007.08297, July 2020.
- [64] Nisha Chandramoorthy and Qiqi Wang. Solenoid map central difference computation. <https://github.com/nishaChandramoorthy/solenoidMap>, 2021.

- [65] Nisha Chandramoorthy and Qiqi Wang. An ergodic-averaging method to differentiate covariant lyapunov vectors. *Nonlinear Dyn*, 104:4083–4102, 2021. doi: 10.1007/s11071-021-06478-0.
- [66] Nisha Chandramoorthy and Qiqi Wang. On the probability of finding non-physical solutions through shadowing. *Journal of Computational Physics*, 440:110389, 2021. ISSN 0021-9991. doi: 10.1016/j.jcp.2021.110389. URL <https://www.sciencedirect.com/science/article/pii/S0021999121002849>.
- [67] Nisha Chandramoorthy and Qiqi Wang. Efficient computation of linear response of chaotic attractors with one-dimensional unstable manifolds. *arXiv e-prints*, art. arXiv:2103.08816, March 2021.
- [68] Nisha Chandramoorthy, Sri Hari Krishna Narayanan, Qiqi Wang, and Paul Hovland. Tangent and adjoint ad-nilss. <https://github.com/nishachandramoorthy/fds>, 2017.
- [69] Nisha Chandramoorthy, Pablo Fernandez, Chaitanya Talnikar, and Qiqi Wang. Feasibility analysis of ensemble sensitivity computation in turbulent flows. *AIAA Journal*, 57(10):4514–4526, 2019. doi: 10.2514/1.J058127. URL <https://doi.org/10.2514/1.J058127>.
- [70] Nisha Chandramoorthy, Zhong-Nan Wang, Qiqi Wang, and Paul Tucker. Toward computing sensitivities of average quantities in turbulent flows. *CoRR*, abs/1902.11112, 2019. URL <http://arxiv.org/abs/1902.11112>.
- [71] Nisha Chandramoorthy, Luca Magri, and Qiqi Wang. Variational optimization and data assimilation in chaotic time-delayed systems with automatic-differentiated shadowing sensitivity. *arXiv e-prints*, art. arXiv:2011.08794, November 2020.
- [72] Mario Chater, Angxiu Ni, Patrick Blonigan, and Qiqi Wang. Least squares shadowing method for sensitivity analysis of differential equations. *SIAM J. Numer. Anal.*, 55(6):3030–3046, 2017. doi: 10.1137/15M1039067.
- [73] Mickaël David Chekroun, J. David Neelin, Dmitri Kondrashov, James C. McWilliams, and Michael Ghil. Rough parameter dependence in climate models and the role of ruelle-pollicott resonances. *Proceedings of the National Academy of Sciences*, 111(5):1684–1690, 2014. ISSN 0027-8424. doi: 10.1073/pnas.1321816111.
- [74] N. I. Chernov. Limit theorems and markov approximations for chaotic dynamical systems. *Probability Theory and Related Fields*, 101(3):321–362, Sep 1995. ISSN 1432-2064. doi: 10.1007/BF01200500. URL <https://doi.org/10.1007/BF01200500>.
- [75] B. T. Chu. On the Energy Transfer to Small Disturbances in Fluid Flow (Part I). *Acta Mechanica*, 1(3):215–234, 1965.

- [76] Vaughn Climenhaga and Anatole Katok. Measure theory through dynamical eyes, 2012.
- [77] Matteo Colangeli and Valerio Lucarini. Elements of a unified framework for response formulae. *Journal of Statistical Mechanics: Theory and Experiment*, 2014(1):P01002, 2014.
- [78] Pierre Collet and J-P Eckmann. Liapunov multipliers and decay of correlations in dynamical systems. *Journal of statistical physics*, 115(1-2):217–254, 2004. doi: 10.1023/B:JOSS.0000019817.71073.61.
- [79] Tim Colonius and Sanjiva K. Lele. Computational aeroacoustics: progress on nonlinear problems of sound generation. *Progress in Aerospace Sciences*, 40(6):345–416, 2004. ISSN 0376-0421. doi: <https://doi.org/10.1016/j.paerosci.2004.09.001>. URL <https://www.sciencedirect.com/science/article/pii/S0376042104000570>.
- [80] Fenwick C Cooper and Peter H Haynes. Climate sensitivity via a nonparametric fluctuation–dissipation theorem. *Journal of the Atmospheric Sciences*, 68(5):937–953, 2011.
- [81] Harry Crimmins and Gary Froyland. Fourier approximation of the statistical properties of anosov maps on tori. *Nonlinearity*, 33(11):6244, 2020. doi: 10.1088/1361-6544/ab987e.
- [82] Dacian N. Daescu, Adrian Sandu, and Gregory R. Carmichael. Direct and adjoint sensitivity analysis of chemical kinetic systems with kpp: Ii—numerical validation and applications. *Atmospheric Environment*, 37(36):5097–5114, 2003. ISSN 1352-2310. doi: <https://doi.org/10.1016/j.atmosenv.2003.08.020>. URL <https://www.sciencedirect.com/science/article/pii/S1352231003006733>.
- [83] Bart de Leeuw, Svetlana Dubinkina, Jason Frank, Andrew Steyer, Xuemin Tu, and Erik Van Vleck. Projected shadowing-based data assimilation. *SIAM J. Appl. Dyn. Syst.*, 17(4):2446—2477, 2018. doi: 10.1137/17M1141163.
- [84] Manfred Denker and Walter Philipp. Approximation by brownian motion for gibbs measures and flows under a function. *Ergodic Theory and Dynamical Systems*, 4(4):541–552, 1984. doi: 10.1017/S0143385700002637.
- [85] Dmitry Dolgopyat. On differentiability of srb states for partially hyperbolic systems. *Inventiones Mathematicae*, 155(2):389–449, 02 2004. URL <https://search.proquest.com/docview/881392378?accountid=12492>. Copyright - Springer-Verlag 2004; Last updated - 2012-06-29.
- [86] M. DØ"RFLE. Spectrum and eigenfunctions of the frobenius-perron operator of the tent map. *Journal of Statistical Physics*, 40(1):93—132, 1985. doi: 10.1007/BF01010528.

- [87] A. P. Dowling. Nonlinear self-excited oscillations of a ducted flame. *Journal of Fluid Mechanics*, 346:271–290, 1997. doi: 10.1017/S0022112097006484. URL http://www.journals.cambridge.org/abstract_S0022112097006484.
- [88] A. P. Dowling. A kinematic model of a ducted flame. *Journal of Fluid Mechanics*, 394:51–72, 1999. doi: 10.1017/S0022112099005686. URL http://www.journals.cambridge.org/abstract_S0022112099005686.
- [89] Semyon Dyatlov. Notes on hyperbolic dynamics, 2018.
- [90] J.-P. Eckmann. Roads to turbulence in dissipative dynamical systems. *Reviews of Modern Physics*, 53(4):643–654, 1981.
- [91] J. P. Eckmann and D. Ruelle. Ergodic theory of chaos and strange attractors. *Reviews of Modern Physics*, 57:617–656, 1985. ISSN 00346861. doi: 10.1103/RevModPhys.57.617.
- [92] Alaa Elmiligui, Khaled Abdol-Hamid, and Edward B Parlette. Detached eddy simulation for the f-16xl aircraft configuration. In *53rd AIAA Aerospace Sciences Meeting*, pages 2015–1496, 2015.
- [93] William Engblom, Abbas Khavaran, and James Bridges. Numerical prediction of chevron nozzle noise reduction using wind-mgbk methodology. In *10th AIAA/CEAS Aeroacoustics Conference*, page 2979, 2004. doi: 10.2514/6.2004-2979.
- [94] Ronald M. Errico. What is an adjoint model? *Bulletin of the American Meteorological Society*, 78(11):2577 – 2592, 01 Nov. 1997. doi: 10.1175/1520-0477(1997)078<2577:WIAAM>2.0.CO;2. URL https://journals.ametsoc.org/view/journals/bams/78/11/1520-0477_1997_078_2577_wiaam_2_0_co_2.xml.
- [95] Ronald M. Errico, Tomislava VukićEvić, and Kevin Raeder. Examination of the accuracy of a tangent linear model. *Tellus A: Dynamic Meteorology and Oceanography*, 45(5):462–477, 1993. doi: 10.3402/tellusa.v45i5.15046. URL <https://doi.org/10.3402/tellusa.v45i5.15046>.
- [96] GL Eyink, TWN Haine, and DJ Lea. Ruelle’s linear response formula, ensemble adjoint schemes and lévy flights. *Nonlinearity*, 17:1867, 2004. doi: 10.1088/0951-7715/17/5/016.
- [97] Florian Feppon and Pierre Lermusiaux. A geometric approach to dynamical model order reduction. *SIAM Journal on Matrix Analysis and Applications*, 39(1):510—538, 2018. doi: 10.1137/16M1095202.
- [98] P Fernandez, NC Nguyen, and J Peraire. Subgrid-scale modeling and implicit numerical dissipation in dg-based large-eddy simulation. In *23rd AIAA Computational Fluid Dynamics Conference*, 2017. doi: 10.2514/6.2017-3951.

- [99] P Fernandez, NC Nguyen, and J Peraire. The hybridized discontinuous galerkin method for implicit large-eddy simulation of transitional turbulent flows. *Journal of Computational Physics*, 336:308–329, 2017. doi: 10.1016/j.jcp.2017.02.015.
- [100] Pablo Fernandez. *Entropy-stable hybridized discontinuous Galerkin methods for large-eddy simulation of transitional and turbulent flows*. PhD thesis, Massachusetts Institute of Technology, 2018.
- [101] Pablo Fernandez and Qiqi Wang. Lyapunov spectrum of the separated flow around the naca 0012 airfoil and its dependence on numerical discretization. *Journal of Computational Physics*, 350:453–469, 2017. doi: 10.1016/j.jcp.2017.08.056.
- [102] A. Fichtner, H.-P. Bunge, and H. Igel. The adjoint method in seismology: I. theory. *Physics of the Earth and Planetary Interiors*, 157(1):86–104, 2006. ISSN 0031-9201. doi: <https://doi.org/10.1016/j.pepi.2006.03.016>. URL <https://www.sciencedirect.com/science/article/pii/S0031920106001051>.
- [103] A. Fichtner, H.-P. Bunge, and H. Igel. The adjoint method in seismology—: II. applications: traveltimes and sensitivity functionals. *Physics of the Earth and Planetary Interiors*, 157(1):105–123, 2006. ISSN 0031-9201. doi: <https://doi.org/10.1016/j.pepi.2006.03.018>. URL <https://www.sciencedirect.com/science/article/pii/S0031920106001038>.
- [104] Krzysztof J Fidkowski and David L Darmofal. Review of output-based error estimation and mesh adaptation in computational fluid dynamics. *AIAA journal*, 49(4):673–694, 2011. doi: 10.2514/1.J050073.
- [105] JB Freund, SK Lele, and P Moin. Numerical simulation of a mach 1.92 turbulent jet and its sound field. *AIAA journal*, 38(11):2023–2031, 2000. doi: 10.2514/2.889.
- [106] S Galatolo and P Giuliatti. A linear response for dynamical systems with additive noise. *Nonlinearity*, 32:2269–2301, 2019. doi: 10.1088/1361-6544/ab0c2e.
- [107] Giovanni Gallavotti. Entropy, thermostats, and chaotic hypothesis. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 16(4):043114, dec 2006. ISSN 1054-1500. doi: 10.1063/1.2372713. URL <http://aip.scitation.org/doi/10.1063/1.2372713>.
- [108] Giovanni Gallavotti and EGD Cohen. Dynamical ensembles in stationary states. *Journal of Statistical Physics*, 80:931–970, 1995. doi: 10.1007/BF02179860.
- [109] K. J. George and R.I. Sujith. Disturbance energy norms: A critical analysis. *Journal of Sound and Vibration*, 331(7):1552–1566, mar 2012. ISSN 0022460X. doi: 10.1016/j.jsv.2011.11.027. URL <http://linkinghub.elsevier.com/retrieve/pii/S0022460X11009059>.

- [110] Michael Ghil and Paola Malanotte-Rizzoli. Data assimilation in meteorology and oceanography. *Advances in geophysics*, 33:141–266, 1991.
- [111] Michael B Giles and Niles A Pierce. An introduction to the adjoint approach to design. *Flow, turbulence and combustion*, 65:393–415, 2000. doi: 10.1023/A:1011430410075.
- [112] Michael B Giles, Mihai C Duta, Jens-Dominik M-uacute, ller, and Niles A Pierce. Algorithm developments for discrete adjoint methods. *AIAA journal*, 41(2):198–205, 2003. doi: 10.2514/2.1961.
- [113] F. Ginelli, P. Poggi, A. Turchi, H. Chaté, R. Livi, and A. Politi. Characterizing dynamics with covariant lyapunov vectors. *Physical Review Letters*, 99(13):130601, 2007. ISSN 00319007. doi: 10.1103/PhysRevLett.99.130601.
- [114] Francesco Ginelli, Hugues Chaté, Roberto Livi, and Antonio Politi. Covariant lyapunov vectors. *Journal of Physics A: Mathematical and Theoretical*, 46(25):254005, jun 2013. doi: 10.1088/1751-8113/46/25/254005.
- [115] Nachiket H Gokhale, Paul E Barbone, and Assad A Oberai. Solution of the nonlinear elasticity imaging inverse problem: the compressible case. *Inverse Problems*, 24(4):045010, 2008. doi: 10.1088/0266-5611/24/4/045010.
- [116] H. Gotoda, H. Nikimoto, T. Miyano, and S. Tachibana. Dynamic properties of combustion instability in a lean premixed gas-turbine combustor. *Chaos*, 21(1):013124, mar 2011. ISSN 1089-7682. doi: 10.1063/1.3563577. URL <http://www.ncbi.nlm.nih.gov/pubmed/21456838>.
- [117] H. Gotoda, T. Ikawa, K. Maki, and T. Miyano. Short-term prediction of dynamical behavior of flame front instability induced by radiative heat loss. *Chaos*, 22(1):033106, sep 2012. ISSN 1089-7682. doi: 10.1063/1.4731267. URL <http://www.ncbi.nlm.nih.gov/pubmed/23020445>.
- [118] Sébastien Gouëzel, Carlangelo Liverani, et al. Compact locally maximal hyperbolic sets for smooth maps: fine statistical properties. *Journal of Differential Geometry*, 79(3):433–477, 2008. doi: 10.4310/jdg/1213798184.
- [119] C. GREBOGI, S. M. HAMMEL, J. A. YORKE, and T. SAUER. Shadowing of physical trajectories in chaotic dynamics: Containment and refinement. *Physical Review Letters*, 65(13):1527–1530, 1990. doi: 10.1103/PhysRevLett.65.1527.
- [120] Celso Grebogi, Edward Ott, and James A. Yorke. Unstable periodic orbits and the dimensions of multifractal chaotic attractors. *Phys. Rev. A*, 37:1711–1724, Mar 1988. doi: 10.1103/PhysRevA.37.1711. URL <https://link.aps.org/doi/10.1103/PhysRevA.37.1711>.
- [121] Andreas Griewank and Andrea Walther. *Evaluating derivatives: principles and techniques of algorithmic differentiation*. SIAM, 2008.

- [122] Andreas Griewank and Andrea Walther. *Jacobian and Hessian Accumulation*, chapter 10, pages 211–243. Society for Industrial and Applied Mathematics, 2008. doi: 10.1137/1.9780898717761.ch10. URL <https://epubs.siam.org/doi/abs/10.1137/1.9780898717761.ch10>.
- [123] Andreas Griewank and Andrea Walther. *Implicit and Iterative Differentiation*, chapter 15, pages 367–396. Society for Industrial and Applied Mathematics, 2008. doi: 10.1137/1.9780898717761.ch15. URL <https://epubs.siam.org/doi/abs/10.1137/1.9780898717761.ch15>.
- [124] Andreas Griewank and Andrea Walther. *Fundamentals of Forward and Reverse*, chapter 3, pages 31–59. Society for Industrial and Applied Mathematics, 2008. doi: 10.1137/1.9780898717761.ch3. URL <https://epubs.siam.org/doi/abs/10.1137/1.9780898717761.ch3>.
- [125] A. Gronskis, D. Heitz, and E. Mémin. Inflow and initial conditions for direct numerical simulation based on adjoint data assimilation. *Journal of Computational Physics*, 242:480–497, 2013. ISSN 0021-9991. doi: <https://doi.org/10.1016/j.jcp.2013.01.051>. URL <https://www.sciencedirect.com/science/article/pii/S0021999113001290>.
- [126] Martin Hairer and Andrew J Majda. A simple framework to justify linear response theory. *Nonlinearity*, 23(4):909–922, mar 2010. doi: 10.1088/0951-7715/23/4/008. URL <https://doi.org/10.1088/0951-7715/23/4/008>.
- [127] S. M. HAMMEL, J. A. YORKE, and C. GREBOGI. Do numerical orbits of chaotic dynamical processes represent true orbits? *Journal of Complexity*, 3(2):136 – 145, 1987. ISSN 0885-064X. doi: 10.1016/0885-064X(87)90024-0.
- [128] Laurent Hascoët and Mauricio Araya-Polo. The adjoint data-flow analyses: Formalization, properties, and applications. *Lecture Notes in Computational Science and Engineering*, 50:135, 2006.
- [129] Laurent Hascoet and Valérie Pascual. The tapenade automatic differentiation tool: principles, model, and specification. *ACM Transactions on Mathematical Software (TOMS)*, 39(3):20, 2013.
- [130] Boris Hasselblatt and Amie Wilkinson. Prevalence of non-Lipschitz Anosov foliations. *Ergodic Theory and Dynamical Systems*, 19(3):643–656, 1999. doi: 10.1017/S0143385799133868. Publisher: Cambridge University Press.
- [131] Sam Hatfield, Andrew McRae, Tim Palmer, and Peter Düben. Single-precision in the tangent-linear and adjoint models of incremental 4d-var. *Monthly Weather Review*, 148(4):1541 – 1552, 01 Apr. 2020. doi: 10.1175/MWR-D-19-0291.1. URL <https://journals.ametsoc.org/view/journals/mwre/148/4/mwr-d-19-0291.1.xml>.

- [132] Guowei He, Guodong Jin, and Yue Yang. Space-time correlations and dynamic coupling in turbulent flows. *Annual Review of Fluid Mechanics*, 49:51–70, 2017. doi: 10.1146/annurev-fluid-010816-060309.
- [133] Patick Heimbach, Dimitris Menemenlis, Martin Losch, Jean-Michel Campin, and Chris Hill. On the formulation of sea-ice models. part 2: Lessons from multi-year adjoint sea-ice export sensitivities through the canadian arctic archipelago. *Ocean Modelling*, 33(1):145–158, 2010.
- [134] Patrick Heimbach and Veronique Bugnion. Greenland ice-sheet volume sensitivity to basal, surface and initial conditions derived from an adjoint model. *Annals of Glaciology*, 50(52):67–80, 2009.
- [135] M. Hénon. *A Two-dimensional Mapping with a Strange Attractor*, pages 94–102. Springer New York, New York, NY, 2004.
- [136] Mark Holland and Ian Melbourne. Central limit theorems and invariance principles for lorenz attractors. *Journal of the London Mathematical Society*, 76: 345–364, 2007. doi: 10.1112/jlms/jdm060.
- [137] Philip Holmes, John L Lumley, Gahl Berkooz, and Clarence W Rowley. *Turbulence, coherent structures, dynamical systems and symmetry*. Cambridge university press, 2012. doi: 10.1017/CBO9780511919701.
- [138] Paul Hovland, Boyana Norris, Michelle Strout, Sanjukta Bhowmick, and Jean Utke. Sensitivity analysis and design optimization through automatic differentiation. In A. Mezzacappa et. al., editor, *SciDAC 2005*, volume 16 of *Journal of Physics: Conference Series*, pages 466–470, Bristol and Philadelphia publisher = Institute of Physics Publishing, 2005. online available at http://www.iop.org/EJ/article/1742-6596/16/1/063/jpconf5_16_063.pdf.
- [139] Jan Hüchelheim, Navjot Kukreja, S. Narayanan, Fabio Luporini, G. Gorman, and P. Hovland. Automatic differentiation for adjoint stencil loops. *Proceedings of the 48th International Conference on Parallel Processing*, 2019.
- [140] Tyler W. Hughes, Momchil Minkov, Ian A. D. Williamson, and Shanhui Fan. Adjoint method and inverse design for nonlinear nanophotonic devices. *ACS Photonics*, 5(12):4781–4787, 2018. doi: 10.1021/acsphotonics.8b01522. URL <https://doi.org/10.1021/acsphotonics.8b01522>.
- [141] Francisco Huhn and Luca Magri. Stability, sensitivity and optimisation of chaotic acoustic oscillations. *Journal of Fluid Mechanics*, 882:A24, 2020. doi: 10.1017/jfm.2019.828. URL <http://arxiv.org/abs/1909.12979>.
- [142] Francisco Huhn and Luca Magri. Optimisation of chaotically perturbed acoustic limit cycles. *Nonlinear Dyn*, 100(2):1641–1657, April 2020. ISSN 0924-090X, 1573-269X. doi: 10.1007/s11071-020-05582-x. URL <http://link.springer.com/10.1007/s11071-020-05582-x>.

- [143] Shaista Hussain, Arindam Basu, Runchun Mark Wang, and Tara Julia Hamilton. Delay learning architectures for memory and classification. *Neurocomputing*, 138:14–26, 2014.
- [144] J.C. Hükelheim, P.D. Hovland, M.M. Strout, and J.-D. Müller. Parallelizable adjoint stencil computations using transposed forward-mode algorithmic differentiation. *Optimization Methods and Software*, 33(4-6):672–693, 2018. doi: 10.1080/10556788.2018.1435654. URL <https://doi.org/10.1080/10556788.2018.1435654>.
- [145] Michael Innes. Don’t Unroll Adjoint: Differentiating SSA-Form Programs. *arXiv e-prints*, art. arXiv:1810.07951, October 2018.
- [146] Mike Innes, Alan Edelman, Keno Fischer, Chris Rackauckas, Elliot Saba, Viral B. Shah, and Will Tebbutt. A differentiable programming system to bridge machine learning and scientific computing. *ArXiv*, abs/1907.07587, 2019.
- [147] S Islam, D Entekhabi, RL Bras, and I Rodriguez-Iturbe. Parameter estimation and sensitivity analysis for the modified bartlett-lewis rectangular pulses model of rainfall. *Journal of Geophysical Research: Atmospheres*, 95(D3):2093–2100, 1990.
- [148] M. Jiang. Differentiating potential functions of srb measures on hyperbolic attractors. *Ergodic Theory and Dynamical Systems*, 32(4):1350–1369, 2012. doi: 10.1017/S0143385711000241.
- [149] Lipika Kabiraj, R I Sujith, and Pankaj Wahi. Bifurcations of Self-Excited Ducted Laminar Premixed Flames. *Journal of Engineering for Gas Turbines and Power*, 134(3):31502, 2011. ISSN 0742-4795. doi: 10.1115/1.4004402. URL <http://dx.doi.org/10.1115/1.4004402>.
- [150] Lipika Kabiraj, Aditya Saurabh, Pankaj Wahi, and R. I. Sujith. Route to chaos for combustion instability in ducted laminar premixed flames. *Chaos*, 22(2): 0–12, 2012. ISSN 10541500. doi: 10.1063/1.4718725.
- [151] F. Kallel and M. Bertrand. Tissue elasticity reconstruction using linear perturbation method. *IEEE Transactions on Medical Imaging*, 15(3):299–313, 1996. doi: 10.1109/42.500139.
- [152] A Karimi and Mark R Paul. Extensive chaos in the lorenz-96 model. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 20:043105, 2010. doi: 10.1063/1.3496397.
- [153] Karthik Kashinath, Iain C. Waugh, and Matthew P. Juniper. Nonlinear self-excited thermoacoustic oscillations of a ducted premixed flame: bifurcations and routes to chaos. *Journal of Fluid Mechanics*, 761:399–430, 2014. ISSN 0022-1120. doi: 10.1017/jfm.2014.601. URL http://www.journals.cambridge.org/abstract_S0022112014006016.

- [154] Anatole Katok and Boris Hasselblatt. *Introduction to the modern theory of dynamical systems*, volume 54. Cambridge university press, 1997.
- [155] Hiroaki Kitano. Systems biology: a brief overview. *Science*, 295(5560):1662–1664, 2002.
- [156] R. Knutti, M. Rugenstein, and G. Hegerl. Beyond equilibrium climate sensitivity. *Nature Geosci*, 10:727–736, 2017. doi: 10.1038/ngeo3017.
- [157] M. KOMURO. Lorenz attractors do not have the pseudo-orbit tracing property. *Journal of the Mathematical Society of Japan*, 37(3):489–514, 1985. doi: 10.2969/jmsj/03730489.
- [158] Yang Kuang. *Delay differential equations: with applications in population dynamics*, volume 191. Academic Press, 1993.
- [159] P. V. Kuptsov and U. Parlitz. Theory and computation of covariant lyapunov vectors. *Journal of Nonlinear Science*, 22:727–762, 2012. ISSN 09388974. doi: 10.1007/s00332-012-9126-5.
- [160] Sergey P. Kuznetsov. Plykin-type attractor in nonautonomous coupled oscillators. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 19(1):013114, 2009. doi: 10.1063/1.3072777. URL <https://doi.org/10.1063/1.3072777>.
- [161] Matieyendou Lamboni, Bertrand Iooss, A-L Popelin, and Fabrice Gamboa. Derivative-based global sensitivity measures: general links with sobol indices and numerical tests. *Mathematics and Computers in Simulation*, 87:45–54, 2013.
- [162] L. D. Landau and E. M. Lifshitz. *Fluid Mechanics*. Pergamon Press, second edition, 1987.
- [163] Lev D Landau. On the problem of turbulence. *Dokl. Akad. Nauk SSSR*, 44(8): 339–349, 1944.
- [164] Rolf H Langland and Nancy L Baker. Estimation of observation impact using the nrl atmospheric variational data assimilation adjoint system. *Tellus A*, 56(3):189–201, 2004.
- [165] D. LASAGNA, A. SHARMA, and J. MEYERS. Periodic shadowing sensitivity analysis of chaotic systems. *Journal of Computational Physics*, 391:119–141, 2019. ISSN 0021-9991. doi: 10.1016/j.jcp.2019.04.021.
- [166] Davide. Lasagna. Sensitivity analysis of chaotic systems using unstable periodic orbits. *SIAM Journal on Applied Dynamical Systems*, 17(1):547–580, 2018. doi: 10.1137/17M114354X. URL <https://doi.org/10.1137/17M114354X>.
- [167] D J Lea, M R Allen, and T W N Haine. Sensitivity analysis of the climate of a chaotic system. *Tellus Series a-Dynamic Meteorology and Oceanography*, 52(5):523–532, 2000. ISSN 1477870X. doi: 10.1256/qj.01.180.

- [168] Daniel J Lea, Myles R Allen, and Thomas WN Haine. Sensitivity analysis of the climate of a chaotic system. *Tellus A: Dynamic Meteorology and Oceanography*, 52:523–532, 2000. doi: 10.1034/j.1600-0870.2000.01137.x.
- [169] F. Ledrappier and L.-S. Young. The metric entropy of diffeomorphisms: Part i: Characterization of measures satisfying pesin’s entropy formula. *Annals of Mathematics*, 122(3):509–539, 1985. ISSN 0003486X. doi: 10.2307/1971328. URL <http://www.jstor.org/stable/1971328>.
- [170] X. LI and S. LIAO. Clean numerical simulation: a new strategy to obtain reliable solutions of chaotic dynamic systems. *Applied Mathematics and Mechanics (English Edition)*, 39(11):1529—1546, 2018. doi: 10.1007/s10483-018-2383-6.
- [171] T. C. Lieuwen and V. Yang. *Combustion Instabilities in Gas Turbine Engines: Operational Experience, Fundamental Mechanisms, and Modeling*. American Institute of Aeronautics and Astronautics, Inc., 2005. ISBN 156347669X.
- [172] Tim Lieuwen. Dynamics of Harmonically Excited Flames. 2013.
- [173] Xu-Dong Liu, Stanley Osher, and Tony Chan. Weighted essentially non-oscillatory schemes. *Journal of computational physics*, 115:200–212, 1994. doi: 10.1006/jcph.1994.1187.
- [174] Carlangelo Liverani. Decay of correlations. *Annals of Mathematics*, 142(2): 239–301, 1995. doi: 10.2307/2118636.
- [175] Carlangelo Liverani. Central limit theorem for deterministic systems. In *International Conference on Dynamical Systems (Montevideo, 1995)*, volume 362, pages 56–75, 1996.
- [176] Carlangelo Liverani. On contact anosov flows. *Annals of Mathematics*, 159 (3):1275–1312, 2004. ISSN 0003486X. URL <http://www.jstor.org/stable/3597178>.
- [177] Brian Lockwood and Dimitri Mavriplis. Gradient-based methods for uncertainty quantification in hypersonic flows. *Computers & Fluids*, 85:27–38, 2013.
- [178] Lord Rayleigh. The explanation of certain acoustical phenomena. *Nature*, 18: 319–321, 1878. URL <http://scholar.google.com/scholar?hl=en&btnG=Search&q=intitle:The+explanation+of+certain+acoustical+phenomena#0>.
- [179] E. N. Lorenz. Deterministic nonperiodic flow. *Journal of Atmospheric Sciences*, 1963. doi: 10.1175/1520-0469(1963)020<0130:DNF>2.0.CO;2.
- [180] Edward N Lorenz. Predictability: A problem partly solved. In *Proc. Seminar on predictability*, volume 1, 1996. doi: 10.1017/CBO9780511617652.004.

- [181] V. Lucarini. Stochastic perturbations to dynamical systems: A response theory approach. *J Stat Phys*, 146:774—786, 2012. doi: 10.1007/s10955-012-0422-0.
- [182] V. Lucarini and S. Sarno. A statistical mechanical approach for the computation of the climatic response to general forcings. *Nonlinear Processes in Geophysics*, 18(1):7–28, 2011. doi: 10.5194/npg-18-7-2011. URL <https://npg.copernicus.org/articles/18/7/2011/>.
- [183] V. Lucarini, R. Blender, C. Herbert, F. Ragone, S. Pascale, and J. Wouters. Mathematical and physical ideas for climate science. *Rev. Geophys.*, 52:809–859, 2014. doi: 10.1002/2013RG000446.
- [184] V. Lucarini, F. Ragone, and F. Lunkeit. Predicting climate change using response theory: Global averages and spatial patterns. *J Stat Phys*, 166:1036–1064, 2017. doi: 10.1007/s10955-016-1506-z.
- [185] Valerio Lucarini. Response operators for markov processes in a finite state space: Radius of convergence and link to the response theory for axiom a systems. *Journal of Statistical Physics*, 162(2):312–333, Jan 2016. ISSN 1572-9613. doi: 10.1007/s10955-015-1409-4. URL <https://doi.org/10.1007/s10955-015-1409-4>.
- [186] Valerio Lucarini and Tamás Bódai. Transitions across melancholia states in a climate model: Reconciling the deterministic and stochastic points of view. *Phys. Rev. Lett.*, 122:158701, Apr 2019. doi: 10.1103/PhysRevLett.122.158701. URL <https://link.aps.org/doi/10.1103/PhysRevLett.122.158701>.
- [187] L. Magri. *Adjoint methods in thermo-acoustic and combustion instability*. PhD thesis, University of Cambridge, 2015.
- [188] L. Magri and M. P. Juniper. Sensitivity analysis of a time-delayed thermo-acoustic system via an adjoint-based approach. *Journal of Fluid Mechanics*, 719:183–202, 2013. doi: 10.1017/jfm.2012.639.
- [189] L. Magri, M. Bauerheim, F. Nicoud, and M. P. Juniper. Stability analysis of thermo-acoustic nonlinear eigenproblems in annular combustors. Part II. Uncertainty quantification. *Journal of Computational Physics*, 325:411–421, 2016. ISSN 10902716. doi: 10.1016/j.jcp.2016.08.043. URL <http://dx.doi.org/10.1016/j.jcp.2016.08.043>.
- [190] Luca Magri. Adjoint Methods as Design Tools in Thermoacoustics. *Applied Mechanics Reviews*, 71(2):020801, 2019. ISSN 00036900. doi: 10.1115/1.4042821.
- [191] Luca Magri and Matthew P Juniper. Sensitivity analysis of a time-delayed thermo-acoustic system via an adjoint-based approach. *Journal of Fluid Mechanics*, 719:183–202, 2013.

- [192] Luca Magri, Michael Bauerheim, Franck Nicoud, and Matthew P. Juniper. Stability analysis of thermo-acoustic nonlinear eigenproblems in annular combustors. part ii. uncertainty quantification. *Journal of Computational Physics*, 325:411–421, 2016. ISSN 0021-9991. doi: <https://doi.org/10.1016/j.jcp.2016.08.043>. URL <https://www.sciencedirect.com/science/article/pii/S0021999116304004>.
- [193] Luca Magri, Matthew P. Juniper, and Jonas P. Moeck. Sensitivity of the rayleigh criterion in thermoacoustics. *Journal of Fluid Mechanics*, 882:R1, 2020. doi: [10.1017/jfm.2019.860](https://doi.org/10.1017/jfm.2019.860).
- [194] Joaquim R. R. A. Martins, Ilan Kroo, and Juan Alonso. An automated method for sensitivity analysis using complex variables. *38th Aerospace Sciences Meeting and Exhibit*, AIAA(2000-0689), 2012. doi: [10.2514/6.2000-689](https://doi.org/10.2514/6.2000-689). URL <https://arc.aiaa.org/doi/abs/10.2514/6.2000-689>.
- [195] Dennis McLaughlin. Recent developments in hydrologic data assimilation. *Reviews of Geophysics*, 33(S2):977–984, 1995. doi: <https://doi.org/10.1029/95RG00740>. URL <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/95RG00740>.
- [196] Antoine McNamara, Adrien Treuille, Zoran Popović, and Jos Stam. Fluid control using the adjoint method. *ACM Trans. Graph.*, 23(3):449–456, August 2004. ISSN 0730-0301. doi: [10.1145/1015706.1015744](https://doi.org/10.1145/1015706.1015744). URL <https://doi.org/10.1145/1015706.1015744>.
- [197] Ian Melbourne and Matthew Nicol. A vector-valued almost sure invariance principle for hyperbolic dynamical systems. *The Annals of Probability*, 37(2):478 – 505, 2009. doi: [10.1214/08-AOP410](https://doi.org/10.1214/08-AOP410). URL <https://doi.org/10.1214/08-AOP410>.
- [198] CFO Mendes, RM da Silva, and MW Beims. Decay of distance autocorrelation and lyapunov exponents. *arXiv preprint arXiv:1903.08202*, 2019.
- [199] G. A. Mensah, L. Magri, and J. P. Moeck. Methods for the calculation of thermoacoustic stability margins and Monte-Carlo free uncertainty quantification. *Journal of Engineering for Gas Turbines and Power*, 140(6):061501, 2018. doi: [10.1115/1.4038156](https://doi.org/10.1115/1.4038156).
- [200] J Miles. Strange attractors in fluid dynamics. *Advances in Applied Mechanics*, 24:189–214, 1984.
- [201] Alan Le Moigne and Ning Qin. Variable-fidelity aerodynamic optimization for turbulent flows using a discrete adjoint formulation. *AIAA journal*, 42(7):1281–1292, 2004. doi: [10.2514/1.2109](https://doi.org/10.2514/1.2109).
- [202] Parviz Moin and John Kim. Numerical investigation of turbulent channel flow. *Journal of fluid mechanics*, 118:341–377, 1982. doi: [10.1017/S0022112082001116](https://doi.org/10.1017/S0022112082001116).

- [203] Siva K. Nadarajah and Antony Jameson. Optimum shape design for unsteady flows with time-accurate continuous and discrete adjoint method. *AIAA Journal*, 45(7):1478–1491, 2007. doi: 10.2514/1.24332. URL <https://doi.org/10.2514/1.24332>.
- [204] V. Nair and R. I. Sujith. A reduced-order model for the onset of combustion instability: Physical mechanisms for intermittency and precursors. *Proceedings of the Combustion Institute, in press*, 35(3):3193–3200, jul 2015. ISSN 15407489. doi: 10.1016/j.proci.2014.07.007. URL <http://linkinghub.elsevier.com/retrieve/pii/S1540748914003174>.
- [205] Vineeth Nair, Gireeshkumaran Thampi, and R I Sujith. Intermittency route to thermoacoustic instability in turbulent combustors. *Journal of Fluid Mechanics*, 756:470–487, 2014. ISSN 1469-7645. doi: 10.1017/jfm.2014.468.
- [206] A. NI and Q. WANG. Sensitivity analysis on chaotic dynamical systems by non-intrusive least squares shadowing (nilss). *Journal of Computational Physics*, 347:56–77, 2017. ISSN 0021-9991. doi: 10.1016/j.jcp.2017.06.033.
- [207] Angxiu Ni. Adjoint shadowing directions in hyperbolic systems for sensitivity analysis, 2018.
- [208] Angxiu Ni. Hyperbolicity, shadowing directions and sensitivity analysis of a turbulent three-dimensional flow. *Journal of Fluid Mechanics*, 863:644–669, 2019. doi: 10.1017/jfm.2018.986.
- [209] Angxiu Ni. Approximating Ruelle’s linear response formula by shadowing methods. *arXiv e-prints*, art. arXiv:2003.09801, March 2020.
- [210] Angxiu Ni. Linear response algorithm for differentiating stationary measures of chaos. *arXiv e-prints*, art. arXiv:2009.00595, September 2020.
- [211] Angxiu Ni and Chaitanya Talnikar. Adjoint sensitivity analysis on chaotic dynamical systems by non-intrusive least squares adjoint shadowing (nilsas). *Journal of Computational Physics*, 395:690 – 709, 2019. ISSN 0021-9991. doi: <https://doi.org/10.1016/j.jcp.2019.06.035>. URL <http://www.sciencedirect.com/science/article/pii/S0021999119304437>.
- [212] Eric J Nielsen and W Kyle Anderson. Aerodynamic design optimization on unstructured meshes using the navier-stokes equations. *AIAA journal*, 37(11): 1411–1419, 1999. doi: 10.2514/2.640.
- [213] Florian Noethen. A projector-based convergence proof of the ginelli algorithm for covariant lyapunov vectors. *Physica D: Nonlinear Phenomena*, 396:18 – 34, 2019. ISSN 0167-2789. doi: <https://doi.org/10.1016/j.physd.2019.02.012>. URL <http://www.sciencedirect.com/science/article/pii/S0167278918302549>.

- [214] a. Orchini, S. J. Illingworth, and M. P. Juniper. Frequency domain and time domain analysis of thermoacoustic oscillations with wave-based acoustics. *Journal of Fluid Mechanics*, 775:387–414, 2015. ISSN 0022-1120. doi: 10.1017/jfm.2015.139. URL http://www.journals.cambridge.org/abstract_S0022112015001391.
- [215] I. Ozaki, F. Kimura, and M. Berz. Higher-order sensitivity analysis of finite element method by automatic differentiation. *Computational Mechanics*, 16: 223–234, 1995. doi: 10.1007/BF00369867.
- [216] Francisco Palacios, Karthik Duraisamy, Juan J Alonso, and Enrique Zuazua. Robust grid adaptation for efficient uncertainty quantification. *AIAA journal*, 50(7):1538–1546, 2012. doi: 10.2514/1.J051379.
- [217] Thomas Sunn Pedersen, Poul K Michelsen, and Jens Juul Rasmussen. Lyapunov exponents and particle dispersion in drift wave turbulence. *Physics of Plasmas*, 3(8):2939–2950, 1996. doi: 10.1063/1.871636.
- [218] Yakov Borisovich Pesin. Characteristic lyapunov exponents and smooth ergodic theory. *Uspekhi Matematicheskikh Nauk*, 32(4):55–112, 1977.
- [219] Jacques EV Peter and Richard P Dwight. Numerical sensitivity analysis for aerodynamic optimization: A survey of approaches. *Computers & Fluids*, 39: 373–391, 2010. doi: 10.1016/j.compfluid.2009.09.013.
- [220] Ben Poole, Subhaneil Lahiri, Maithra Raghu, Jascha Sohl-Dickstein, and Surya Ganguli. Exponential expressivity in deep neural networks through transient chaos. In D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016. URL <https://proceedings.neurips.cc/paper/2016/file/148510031349642de5ca0c544f31b2ef-Paper.pdf>.
- [221] C. Pugh and M. Shub. Stable ergodicity. *Bulletin of the American Mathematical Society (N.S.)*, 41:1–41, 2004. doi: 10.1090/S0273-0979-03-00998-4.
- [222] Thomas H Pulliam and John A Vastano. Transition to Chaos in an Open Unforced 2D Flow. *Journal of Computational Physics*, 105:133–149, 1993. doi: 10.1006/jcph.1993.1059.
- [223] C. Quinn, J. Sieber, and A. von der Heydt. Effects of periodic forcing on a paleoclimate delay model. *SIAM J. Appl. Dyn. Syst.*, 18(2):1060–1077, 2019. doi: 10.1137/18M1203079.
- [224] Joaquim R RA Martins, Juan J Alonso, and James J Reuther. High-fidelity aerostructural design optimization of a supersonic business jet. *Journal of Aircraft*, 41(3):523–530, 2004. doi: 10.2514/1.11478.

- [225] Chris Rackauckas, Mike Innes, Yingbo Ma, Jesse Bettencourt, Lyndon White, and Vaibhav Dixit. *Diffeqflux.jl*-a julia library for neural differential equations. *arXiv preprint arXiv:1902.02376*, 2019.
- [226] Christopher Rackauckas and Qing Nie. *Differentialequations.jl*-a performant and feature-rich ecosystem for solving differential equations in julia. *Journal of Open Research Software*, 5(1), 2017.
- [227] Christopher Rackauckas, Yingbo Ma, Vaibhav Dixit, Xingjian Guo, Mike Innes, Jarrett Revels, Joakim Nyberg, and Vijay D Ivaturi. A comparison of automatic differentiation and continuous sensitivity analysis for derivatives of differential equation solutions. *ArXiv*, abs/1812.01892, 2018.
- [228] Francesco Ragone, Valerio Lucarini, and Frank Lunkeit. A new framework for climate sensitivity and prediction: a modelling perspective. *Climate Dynamics*, 46:1459–1471, 2016. doi: 10.1007/s00382-015-2657-3.
- [229] Donald P Rizzetta and Miguel R Visbal. Large-eddy simulation of supersonic cavity flowfields including flow control. *AIAA journal*, 41(8):1452–1462, 2003. doi: 10.2514/2.2128.
- [230] Stephen K Robinson. Coherent motions in the turbulent boundary layer. *Annual Review of Fluid Mechanics*, 23(1):601–639, 1991. doi: 10.1146/annurev.fl.23.010191.003125.
- [231] Gerard H. Roe and Marcia B. Baker. Why is climate sensitivity so unpredictable? *Science*, 318(5850):629–632, 2007. doi: 10.1126/science.1144735.
- [232] David Ruelle. Differentiation of srb states. *Communications in Mathematical Physics*, 187:227–241, 1997. doi: 10.1007/s002200050134.
- [233] David Ruelle. General linear response formula in statistical mechanics, and the fluctuation-dissipation theorem far from equilibrium. *Physics Letters A*, 245: 220–224, 1998. doi: 10.1016/S0375-9601(98)00419-8.
- [234] David Ruelle. Differentiation of srb states: correction and complements. *Communications in mathematical physics*, 234:185–190, 2003. doi: 10.1007/s00220-002-0779-z.
- [235] David Ruelle. Differentiation of srb states for hyperbolic flows. *Ergodic Theory and Dynamical Systems*, 28:613–631, 2008. doi: 10.1017/S0143385707000260.
- [236] David Ruelle. A review of linear response theory for general differentiable dynamical systems. *Nonlinearity*, 22:855, 2009. doi: 10.1088/0951-7715/22/4/009.
- [237] Andrea Saltelli, Marco Ratto, Terry Andres, Francesca Campolongo, Jessica Cariboni, Debora Gatelli, Michaela Saisana, and Stefano Tarantola. *Global sensitivity analysis: the primer*. John Wiley & Sons, 2008.

- [238] Barry Saltzman. Finite amplitude free convection as an initial value problem—i. *Journal of Atmospheric Sciences*, 19(4):329 – 341, 1962. doi: 10.1175/1520-0469(1962)019<0329:FAFCAA>2.0.CO;2. URL https://journals.ametsoc.org/view/journals/atasc/19/4/1520-0469_1962_019_0329_fafcaa_2_0_co_2.xml.
- [239] Adrian Sandu, Dacian N Daescu, and Gregory R Carmichael. Direct and adjoint sensitivity analysis of chemical kinetic systems with kpp: Part i—theory and software tools. *Atmospheric Environment*, 37(36):5083–5096, 2003.
- [240] T. SAUER. Computer arithmetic and sensitivity of natural measure. *Journal of Difference Equations and Applications*, 11(7):669–676, 2005. doi: 10.1080/10236190412331334545.
- [241] Oliver T Schmidt, Aaron Towne, Georgios Rigas, Tim Colonius, and Guillaume A Brès. Spectral analysis of jet turbulence. *Journal of Fluid Mechanics*, 855:953–982, 2018. doi: 10.1017/jfm.2018.675.
- [242] S. C. Sherwood, M. J. Webb, J. D. Annan, K. C. Armour, P. M. Forster, J. C. Hargreaves, G. Hegerl, S. A. Klein, K. D. Marvel, E. J. Rohling, M. Watanabe, T. Andrews, P. Braconnot, C. S. Bretherton, G. L. Foster, Z. Hausfather, A. S. von der Heydt, R. Knutti, T. Mauritsen, J. R. Norris, C. Proistosescu, M. Rugenstein, G. A. Schmidt, K. B. Tokarska, and M. D. Zelinka. An assessment of earth’s climate sensitivity using multiple lines of evidence. *Reviews of Geophysics*, 58(4):e2019RG000678, 2020. doi: <https://doi.org/10.1029/2019RG000678>. URL <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2019RG000678>. e2019RG000678 2019RG000678.
- [243] C. F. Silva, L. Magri, T. Runte, and W. Polifke. Uncertainty Quantification of Growth Rates of Thermoacoustic Instability by an Adjoint Helmholtz Solver. *Journal of Engineering for Gas Turbines and Power*, 139(1):011901, 2016. ISSN 0742-4795. doi: 10.1115/1.4034203. URL <http://gasturbinespower.asmedigitalcollection.asme.org/article.aspx?doi=10.1115/1.4034203>.
- [244] Lawrence Sirovich and Anil E Deane. A computational study of rayleigh–bénard convection. part 2. dimension considerations. *Journal of fluid mechanics*, 222: 251–265, 1991. doi: 10.1017/S002211209100109X.
- [245] J Slipantschuk, O F Bandtlow, and W Just. Complete spectral data for analytic Anosov maps of the torus. *Nonlinearity*, 30(7):2667–2686, July 2017. ISSN 0951-7715, 1361-6544. doi: 10.1088/1361-6544/aa700f. URL <https://iopscience.iop.org/article/10.1088/1361-6544/aa700f>.
- [246] Julia Slipantschuk, Oscar F Bandtlow, and Wolfram Just. On the relation between lyapunov exponents and exponential decay of correlations. *Journal of Physics A: Mathematical and Theoretical*, 46(7):075101, 2013. doi: 10.1088/1751-8113/46/7/075101.

- [247] Adam A. Sliwiak, Nisha Chandramoorthy, and Qiqi Wang. Ergodic Sensitivity Analysis of One-Dimensional Chaotic Maps. *arXiv e-prints*, art. arXiv:2007.15594, July 2020.
- [248] Adam A. Sliwiak, Nisha Chandramoorthy, and Qiqi Wang. Computational assessment of smooth and rough parameter dependence of statistics in chaotic dynamical systems. *arXiv e-prints*, art. arXiv:2101.08379, January 2021.
- [249] Hal L Smith. *An introduction to delay differential equations with applications to the life sciences*, volume 57. Springer New York, 2011.
- [250] Jaroslaw Sobieszczanski-Sobieski. Sensitivity analysis and multidisciplinary optimization for aircraftdesign-recent advances and results. *Journal of Aircraft*, 27(12):993–1001, 1990.
- [251] I.M. Sobol and S. Kucherenko. Derivative based global sensitivity measures and their link with global sensitivity indices. *Mathematics and Computers in Simulation*, 79(10):3009–3017, 2009.
- [252] Kyle D Squires, James R Forsythe, Scott A Morton, Douglas C Blake, Matthieu Serrano, Ken E Wurtzler, William Z Strang, Robert F Tomaro, and PR Spalart. Analysis of full aircraft with massive separation using detached-eddy simulation. In *Proceedings of the High Performance Computing Modernization Program 2002 Users Group Conference, Austin, TX, June*, pages 9–14, 2002.
- [253] Priya Subramanian, RI Sujith, and P Wahi. Subcritical bifurcation and bistability in thermoacoustic systems. *Journal of Fluid Mechanics*, 715:210–238, 2013.
- [254] N.-Z. Sun and W. W.-G. Yeh. Coupled inverse problems in groundwater modeling: 2. identifiability and experimental design. *Water Resour. Res.*, 26(10):2527—2540, 1990. doi: 10.1029/WR026i010p02527.
- [255] W. W.-G. Sun, N.-Z.and Yeh. Coupled inverse problems in groundwater modeling: 1. sensitivity analysis and parameter identification. *Water Resour. Res.*, 26(10):2507—2525, 1990. doi: 10.1029/WR026i010p02507.
- [256] Olivier Talagrand and Philippe Courtier. Variational assimilation of meteorological observations with the adjoint vorticity equation. i: Theory. *Quarterly Journal of the Royal Meteorological Society*, 113(478):1311–1328, 1987.
- [257] Chaitanya Talnikar, Qiqi Wang, and Gregory M Laskowski. Unsteady adjoint of pressure loss for a fundamental transonic turbine vane. *Journal of Turbomachinery*, 139:031001, 2017. doi: 10.1115/1.4034800.
- [258] Albert Tarantola. *Inverse problem theory and methods for model parameter estimation*. SIAM, 2005.

- [259] T. Traverso, A. Bottaro, and L. Magri. Data assimilation in thermoacoustic instability with Lagrangian optimization. *EuroMech Vienna*, 2018.
- [260] Tullio Traverso and Luca Magri. Data Assimilation in a Nonlinear Time-Delayed Dynamical System with Lagrangian Optimization. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, volume 11539 LNCS, pages 156–168. Springer, 2019. ISBN 9783030227463. doi: 10.1007/978-3-030-22747.
- [261] L. N. Trefethen and M. Embree. *Spectra and Pseudospectra*. Princeton University Press, 2005.
- [262] A. Trevisan, M. D’Isidoro, and O. Talagrand. Four-dimensional variational assimilation in the unstable subspace (4DVar-AUS) and the optimal subspace dimension. *Quarterly Journal of the Royal ...*, 133(October):937–948, 2007. ISSN 00359009. doi: 10.1002/qj. URL <http://onlinelibrary.wiley.com/doi/10.1002/qj.71/abstract>.
- [263] Paul G Tucker. Novel miles computations for jet flows and noise. *International Journal of Heat and Fluid Flow*, 25(4):625–635, 2004. doi: 10.1016/j.ijheatfluidflow.2003.11.021.
- [264] James C. Tyacke and Paul G. Tucker. Future Use of Large Eddy Simulation in Aero-engines. *Journal of Turbomachinery*, 137(8), 08 2015. ISSN 0889-504X. doi: 10.1115/1.4029363. URL <https://doi.org/10.1115/1.4029363>. 081005.
- [265] Jean Utke and Uwe Naumann. Combinatorial problems in Open/AD. In Uwe Naumann and Olaf Schenk, editors, *Combinatorial Scientific Computing*, Computational Science. Chapman and Hall/CRC, January 2012. chap. 6.
- [266] Jean Utke, Uwe Naumann, Mike Fagan, Nathan Tallent, Michelle Strout, Patrick Heimbach, Chris Hill, and Carl Wunsch. Openad/f: A modular open-source tool for automatic differentiation of fortran codes. *ACM Transactions on Mathematical Software (TOMS)*, 34(4):18, 2008.
- [267] Daniele Venturi, Heyrim Cho, and George Em Karniadakis. Mori-zwanzig approach to uncertainty quantification. *Handbook of Uncertainty Quantification*. Springer, 2016. doi: 10.1007/978-3-319-11259-6_28-2.
- [268] Felipe AC Viana, Timothy W Simpson, Vladimir Balabanov, and Vasilli Toropov. Special section on multidisciplinary design optimization: Metamodeling in multidisciplinary design optimization: How far have we really come? *AIAA Journal*, 2014.
- [269] H.X. Vo and L.J. Durlofsky. Data assimilation and uncertainty assessment for complex geological models using a new pca-based parameterization. *Comput Geosci*, 19:747—767, 2015. doi: 10.1007/s10596-015-9483-x.

- [270] Qiqi Wang. *Uncertainty quantification for unsteady fluid flow using adjoint-based approaches*. PhD thesis, Stanford University, 2008.
- [271] Qiqi Wang. Convergence of the least squares shadowing method for computing derivative of ergodic averages. *SIAM Journal on Numerical Analysis*, 52:156–170, 2014. doi: 10.1137/130917065.
- [272] Qiqi Wang, Karthik Duraisamy, Juan J Alonso, and Gianluca Iaccarino. Risk assessment of scramjet unstart using adjoint-based sampling methods. *AIAA journal*, 50(3):581–592, 2012. doi: 10.2514/1.J051264.
- [273] Qiqi Wang, Rui Hu, and Patrick Blonigan. Least squares shadowing sensitivity analysis of chaotic limit cycle oscillations. *Journal of Computational Physics*, 267:210 – 224, 2014. ISSN 0021-9991. doi: 10.1016/j.jcp.2014.03.002. URL <http://www.sciencedirect.com/science/article/pii/S0021999114001715>.
- [274] I. C. Waugh, K. Kashinath, and M. P. Juniper. Matrix-free continuation of limit cycles and their bifurcations for a ducted premixed flame. *Journal of Fluid Mechanics*, 759:1–27, 2014. ISSN 0022-1120. doi: 10.1017/jfm.2014.549. URL http://www.journals.cambridge.org/abstract_S0022112014005497.
- [275] C. Wormell. Spectral galerkin methods for transfer operators in uniformly expanding dynamics. *Numer. Math.*, 142:421—463, 2019. doi: 10.1007/s00211-019-01031-z.
- [276] Caroline L. Wormell and Georg A. Gottwald. On the validity of linear response theory in high-dimensional deterministic dynamical systems. *J Stat Phys*, 172:1479—1498, 2018. doi: 10.1007/s10955-018-2106-x.
- [277] Caroline L. Wormell and Georg A. Gottwald. Linear response for macroscopic observables in high-dimensional systems. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 29(11):113127, 2019. doi: 10.1063/1.5122740. URL <https://doi.org/10.1063/1.5122740>.
- [278] Carl Wunsch, Patrick Heimbach, Rui M Ponte, and Ichiro Fukumori. The global general circulation of the ocean estimated by the ecco-consortium. *Oceanography*, 22(2):88–103, 2009.
- [279] T. YOSHIDA, H. MORI, and H. SHIGEMATSU. Analytic study of chaos of the tent map: Band structures, power spectra, and critical behaviors. *Journal of Statistical Physics*, 31(2):279–308, 1983. doi: 10.1007/BF01011583.
- [280] Lai-Sang Young. Statistical properties of dynamical systems with some hyperbolicity. *Annals of Mathematics*, 147:585–650, 1998. doi: 10.2307/120960.
- [281] Lai-Sang Young. What are srb measures, and which dynamical systems have them? *Journal of Statistical Physics*, 108:733–754, 2002. doi: 10.1023/A:1019762724717.

- [282] Mark D. Zelinka, Timothy A. Myers, Daniel T. McCoy, Stephen Po-Chedley, Peter M. Caldwell, Paulo Ceppi, Stephen A. Klein, and Karl E. Taylor. Causes of higher climate sensitivity in cmip6 models. *Geophysical Research Letters*, 47(1):e2019GL085782, 2020. doi: <https://doi.org/10.1029/2019GL085782>. URL <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2019GL085782>. e2019GL085782 10.1029/2019GL085782.
- [283] Lin Zhang, E.M. Constantinescu, A. Sandu, Y. Tang, T. Chai, G.R. Carmichael, D. Byun, and E. Olaguer. An adjoint sensitivity analysis and 4d-var data assimilation study of texas air quality. *Atmospheric Environment*, 42(23): 5787–5804, 2008. ISSN 1352-2310. doi: <https://doi.org/10.1016/j.atmosenv.2008.03.048>. URL <https://www.sciencedirect.com/science/article/pii/S1352231008003178>. Selected Papers from the First International Conference on Atmospheric Chemical Mechanisms.
- [284] Wei Zhang, Michael S Branicky, and Stephen M Phillips. Stability of networked control systems. *IEEE Control Systems*, 21(1):84–99, 2001.