Quantum algorithms for testing probability distributions

Sergey Bravyi[1] Aram Harrow[2] Avinatan Hassidim[3]

[1]IBM

[2]University of Bristol

[3]MIT

STACS
4 March, 2010

The problem: testing probability distributions

- We are given samples of $[N]$ drawn according to p.
- The goal is to determine some (often symmetric) property of p. For example:
  - Entropy: $H(p) = \sum_{i=1}^{N} p_i \log \frac{1}{p_i}$.
  - Distance from uniform distribution:

$$\frac{1}{2}\|p-u\| = \frac{1}{2}\sum_{i=1}^{N}\left|p_i - \frac{1}{N}\right|$$

- Alternatively, we can draw samples according to p or q.
  - Statistical distance: Given thresholds $0 \le \alpha < \beta \le 1$, determine whether $\frac{1}{2}\|p-q\|_1 \le \alpha$ or $\frac{1}{2}\|p-q\|_1 \ge \beta$.
  - Special cases include when $\alpha = 0$ (determining whether $p = q$) and when $\beta = 1$ (determining whether p and q have orthogonal support.)

Motivation

- This is a basic example of property testing.
- It is a primitive in other tasks, such as testing whether a graph is bipartite.
- Estimating trace distance is a complete problem for SZK. Quantum computers can speedup the naive algorithm for solving problems in NP. Can they do the same for SZK?

What it means to sample on a quantum computer

There is no canonical answer. Let $p \in \mathbb{R}^N$ be a probability distribution. Here are three possibilities, in order of increasing strength.

	Model	Cost of uniformity testing
1	The ability to prepare an $N \times N$ density matrix ρ with spec $\rho = p$	$\Theta(N)$
2	The existence of an efficient classical circuit that can sample from p.	$\sqrt{N}$ classically, $N^{1/3}$ quantumly
3	The ability to prepare $\sum_{i=1}^{N} \sqrt{p_i}\, \|i\rangle$.	$O(1)$

Scenario 1 is weaker and scenario 3 is stronger.

In this talk, we will focus on scenario 2.

Defining sampling oracles

We want to treat the classical algorithm creating samples of p as a black box.

$\boxed{\text{Sampling Algorithm}} \longrightarrow i$ with probability p_i

However, this model is too restrictive.

Random seed $r \in \{0,1\}^m \longrightarrow \boxed{\text{Sampling Algorithm}} \longrightarrow i = f(r)$ where $p_i = \frac{|f^{-1}(i)|}{2^m}$

Our quantum algorithm will make use of oracle access to f.

Classical results

The symmetric case was mostly solved by [Valiant, STOC '08].

Canonical tester

1. Draw M samples according to p.
2. Suppose that item i appears $s(i)$ times.
3. If $s(i) \geq \theta$, then estimate $\hat{p}_i = s(i)/M$. Otherwise, consider the range $\hat{p}_i \in [0, \frac{\theta}{M}]$.
4. Hope that $\hat{p}$ gives an unambiguous answer.

Applications

- Estimating trace distance in general requires $N^{1-o(1)}$ samples.
- Determining whether $p = q$ or $\frac{1}{2}\|p - q\|_1 \geq \epsilon$ requires $\Theta(N^{2/3})$ samples.

Previous quantum results

- The first example of quantum advantage is Grover's 1996 search algorithm. (Proved optimal by BBBV in 1994!)
- For any subset $S \subset [N]$, let $\pi = \sum_{i \in S} p_i$. Grover's algorithm can determine whether $\pi = 0$ or $\pi \geq \theta$ in time $O(1/\sqrt{\theta})$.
- This can be used distinguish 1-1 functions from 2-1 functions in time $O(N^{1/3})$. [Brassard, Høyer, Tapp; quant-ph/9705002]
- More generally, we can output $\pi \pm O(\epsilon)$ in time $O(\sqrt{\pi}/\epsilon)$. [Brassard, Høyer, Mosca, Tapp; quant-ph/0005055] Compare with $O(\pi/\epsilon^2)$ for classical sampling.
- [Aaronson and Ambainis; arXiv:0911.0996] prove that for symmetric problems, any Q-query quantum algorithm can be turned into an $O(Q^9)$-query randomized algorithm.

Our results

Given two distributions p, q and constants $0 < \epsilon \leq \theta \leq 1$ we consider three problems.

Goal	to distinguish	
Uniformity testing	$p = u$	$\frac{1}{2}\|p - u\|_1 \geq \epsilon$
Statistical distance	$\frac{1}{2}\|p - q\|_1 \leq \theta - \epsilon$	$\frac{1}{2}\|p - q\|_1 \geq \theta$
Orthogonality	$\frac{1}{2}\|p - q\|_1 \leq 1 - \epsilon$	$\frac{1}{2}\|p - q\|_1 = 1$

(u denotes the uniform distribution on $[N]$.)

Results:

Goal	Classical	Quantum
Uniformity testing	$\Theta(N^{1/2})$	$\Theta(N^{1/3})$
Statistical distance	$N^{1-o(1)}$	$O(N^{1/2})$
Orthogonality	$\Theta(N^{1/2})$	$\Theta(N^{1/3})$

(Uniformity lower bound from [Chakraborty *et. al*, unpublished].)

Distribution testing protocols

Algorithm for statistical distance

Consider the r.v. X which equals $\frac{|p_i - q_i|}{p_i + q_i}$ with probability $\frac{1}{2}(p_i + q_i)$.

- $\mathbb{E}(X) = \frac{1}{2}\|p - q\|_1$
- $\text{Var}(X) \leq \mathbb{E}(X^2) \leq 1$
- Estimating X to constant multiplicative accuracy requires $O(\sqrt{N/\delta})$ quantum queries when $\max(p_i, q_i) \geq \delta/2N$. This happens with probability $\geq 1 - \delta$.
- Therefore $O(\sqrt{N})$ queries suffice.

Algorithm for fidelity: $\sum_{i=1}^N \sqrt{p_i q_i}$

Now let X equal $\sqrt{p_i/q_i}$ with probability q_i.

- $\mathbb{E}(X) = \sum_{i=1}^N \sqrt{p_i q_i}$ $\quad$ $\text{Var}(X) \leq \mathbb{E}(X^2) = \sum_{i=1}^N p_i = 1$
- Again $O(\sqrt{N})$ queries suffice.

Distribution testing protocols

Algorithm for uniformity testing

- Take $M \sim N^{1/3}$.
- Sample $S = \{i_1, \ldots, i_M\}$ according to p.
- If there is a collision, then output "not uniform."
- Let $p_S = p_{i_1} + \ldots + p_{i_M}$.
- Estimate p_S to constant accuracy and output "uniform" iff $p_S \approx M/N$.

Algorithm for orthogonality testing

- Take $M \sim N^{1/3}$.
- Sample $S = (i_1, \ldots, i_M)$ according to p, ignoring duplicates.
- Estimate $q_S = q_{i_1} + \ldots + q_{i_M}$ and output "orthogonal" iff the estimate is 0.

Discussion

Unlike the classical "canonical tester," there are many different quantum approaches.

It is unknown whether there is a general framework that encompasses all optimal quantum algorithms for testing symmetric properties of distributions.

The only general purpose lower bound is the Aaronson-Ambainis result. It is probably not tight, and does not suggest a canonical algorithm.

One other subtlety: what if we have a quantum algorithm that can generate samples according to p? Now there is no seed, but the subroutine to estimate probabilities still works. Is this ever a weaker primitive?