


Learning the closest product state

John Bostanci

Joint work with Ainesh Bakshi, William Kretschmer, Ryan O'Donnell,
Zeph Landau, Jerry Li, Allen Liu, and Ewin Tang


Motivation

Let's imagine the following scenario:




Motivation

Let's imagine the following scenario:




Motivation

Let's imagine the following scenario:



Motivation

Let's imagine the following scenario:



Agnostic quantum tomography

Given a model class C and copies of an arbitrary quantum state ρ , output a description of the closest state in C to ρ .

$$(C, \rho^{\otimes n}) \rightarrow \operatorname{argmax}_{|\psi\rangle \in C} \langle \psi | \rho | \psi \rangle$$

Agnostic quantum tomography

Given a model class C and copies of an arbitrary quantum state ρ , output a description of the closest state in C to ρ .

$$(C, \rho^{\otimes n}) \rightarrow \operatorname{argmax}_{|\psi\rangle \in C} \langle \psi | \rho | \psi \rangle$$


OPT

Agnostic quantum tomography

Given a model class $\mathcal{C}$ and copies of an arbitrary quantum state ρ , output a description of the closest state in $\mathcal{C}$ to ρ .

$$(\mathcal{C}, \rho^{\otimes n}) \rightarrow |\phi\rangle \in \mathcal{C} : \text{OPT} - \langle \phi | \rho | \phi \rangle \leq \epsilon$$

Agnostic quantum tomography

Given a model class C and copies of an arbitrary quantum state ρ , output a description of the closest state in C to ρ .

$$(C, \rho^{\otimes n}) \rightarrow |\phi\rangle \in C: \text{OPT} - \langle \phi | \rho | \phi \rangle \leq \epsilon$$

Shadow tomography provides a sample efficient, but not computationally efficient, solution.

Agnostic tomography is hard

Even for “simple” model classes like product states, agnostic tomography isn’t immediate.

Agnostic tomography is hard

Even for “simple” model classes like product states, agnostic tomography isn’t immediate. Consider the following state:


$$\sqrt{1 - \epsilon} |0^n\rangle + \sqrt{\epsilon} |+\rangle^n$$

Agnostic tomography is hard

Even for “simple” model classes like product states, agnostic tomography isn’t immediate. Consider the following state:

$$\sqrt{1 - \epsilon} |0^n\rangle + \sqrt{\epsilon} |+\rangle^n$$

The closest product state is $|0^n\rangle$




Agnostic tomography is hard

Even for “simple” model classes like product states, agnostic tomography isn’t immediate. Consider the following state:

$$\sqrt{1 - \epsilon} |0^n\rangle + \sqrt{\epsilon} |+\rangle^n$$

The closest product state is $|0^n\rangle$



Every marginal is $\approx \sqrt{1 - \epsilon} |0\rangle + \sqrt{\epsilon} |1\rangle$.
Taking the tensor product is bad!

Our results (1)

Main result: there is an algorithm for agnostic tomography of product states that has sample and time complexity

$$\text{poly}(n^{\text{poly}(\frac{1}{\epsilon})})$$

Our results (2)

Suppose there was an polynomial-time algorithm for agnostic tomography of product states when ϵ is inverse polynomial in n , i.e. outputting a state satisfying

$$\langle \phi | \rho | \phi \rangle \geq \text{OPT} - \frac{1}{\text{poly}(n)}$$

Then BQP contains NP.

Our results (3)

In some settings, we can get fully polynomial run-times!

- When $\text{OPT} \geq \frac{5}{6}$, we give an algorithm that uses $O\left(\frac{n}{\epsilon}\right)$ copies and runs in time $O\left(\frac{n^2}{\epsilon} \log n\right)$.

Our results (3)

In some settings, we can get fully polynomial run-times!

- When $\text{OPT} \geq \frac{5}{6}$, we give an algorithm that uses $O\left(\frac{n}{\epsilon}\right)$ copies and runs in time $O\left(\frac{n^2}{\epsilon} \log n\right)$.
- When we have the promise that each qudit can only be one of s states, and all of the states have fidelity at most $1 - \delta$, we give an algorithm that runs in time $\text{poly}(ns)^{\log(1/\epsilon)/\delta}$.

Our results (3)

In some settings, we can get fully polynomial run-times!

- When $\text{OPT} \geq \frac{5}{6}$, we give an algorithm that uses $O\left(\frac{n}{\epsilon}\right)$ copies and runs in time $O\left(\frac{n^2}{\epsilon} \log n\right)$.
- When we have the promise that each qudit can only be one of s states, and all of the states have fidelity at most $1 - \delta$, we give an algorithm that runs in time $\text{poly}(ns)^{\log(1/\epsilon)/\delta}$.
- We give an improper tomography algorithm for matrix product states that runs in time $\text{poly}(n, 1/\epsilon)$.

Rest of the talk

The goal of the rest of the talk is to give an overview and intuition for the main algorithm.

- Going from local to global

Rest of the talk

The goal of the rest of the talk is to give an overview and intuition for the main algorithm.

- Going from local to global
- Reduction to polynomial optimization

Rest of the talk

The goal of the rest of the talk is to give an overview and intuition for the main algorithm.

- Going from local to global
- Reduction to polynomial optimization
- Solving the polynomial optimization

Rest of the talk

The goal of the rest of the talk is to give an overview and intuition for the main algorithm.

- Going from local to global
- Reduction to polynomial optimization
- Solving the polynomial optimization
- Putting everything together

Going from local to global

We saw before that simply taking the single qubit marginals of the state and taking their tensor product is a bad idea.

Going from local to global

We saw before that simply taking the single qubit marginals of the state and taking their tensor product is a bad idea.

We can refine our approach a little bit!

First strategy: Let $|\pi\rangle$ have $\text{OPT} - \epsilon$ fidelity with the first $k-1$ qubits. Take an ϵ -net over states $\{|\pi_k\rangle\}$, and test if any of the following states have high fidelity ρ :

$$|\pi\rangle \otimes |\pi_k\rangle$$

First strategy: Let $|\pi\rangle$ have $\text{OPT} - \epsilon$ fidelity with the first $k-1$ qubits. Take an ϵ -net over states $\{|\pi_k\rangle\}$, and test if any of the following states have high fidelity ρ :

$$|\pi\rangle \otimes |\pi_k\rangle$$

However: If $|\pi\rangle$ has fidelity $\text{OPT} - \epsilon$, and we're taking an ϵ -net, every state of this form could have fidelity $\text{OPT} - 2\epsilon$.

First strategy: Let $|\pi\rangle$ have $\text{OPT} - \epsilon$ fidelity with the first $k-1$ qubits. Take an ϵ -net over states $\{|\pi_k\rangle\}$, and test if any of the following states have high fidelity ρ :


$$|\pi\rangle \otimes |\pi_k\rangle$$

However: If $|\pi\rangle$ has fidelity $\text{OPT} - \epsilon$, and we're taking an ϵ -net, every state of this form could have fidelity $\text{OPT} - 2\epsilon$.

We lose a little on every step!
Can we fix this loss?

Problem: Let $|\pi\rangle$ be a product state such that $\langle\pi|\rho|\pi\rangle \geq \text{OPT} - 10\epsilon$,
can we find a nearby $|\pi'\rangle$ such that $\langle\pi'|\rho|\pi'\rangle \geq \text{OPT} - \epsilon$?

We found $|\pi\rangle$ via a **local** update (appending $|\pi_k\rangle$)



Problem: Let $|\pi\rangle$ be a product state such that $\langle\pi|\rho|\pi\rangle \geq \text{OPT} - 10\epsilon$,
can we find a nearby $|\pi'\rangle$ such that $\langle\pi'|\rho|\pi'\rangle \geq \text{OPT} - \epsilon$?

We found $|\pi\rangle$ via a **local** update (appending $|\pi_k\rangle$)

Problem: Let $|\pi\rangle$ be a product state such that $\langle\pi|\rho|\pi\rangle \geq \text{OPT} - 10\epsilon$,
can we find a nearby $|\pi'\rangle$ such that $\langle\pi'|\rho|\pi'\rangle \geq \text{OPT} - \epsilon$?

$|\pi'\rangle$ should incorporate something **global** around ρ

We found $|\pi\rangle$ via a **local** update (appending $|\pi_k\rangle$)

Problem: Let $|\pi\rangle$ be a product state such that $\langle\pi|\rho|\pi\rangle \geq \text{OPT} - 10\epsilon$,
can we find a nearby $|\pi'\rangle$ such that $\langle\pi'|\rho|\pi'\rangle \geq \text{OPT} - \epsilon$?

$|\pi'\rangle$ should incorporate something **global** around ρ

Searching only near $|\pi\rangle$ makes things **efficient**

Problem: Let $|\pi\rangle$ be a product state such that $\langle\pi|\rho|\pi\rangle \geq \text{OPT} - 10\epsilon$,
can we find a nearby $|\pi'\rangle$ such that $\langle\pi'|\rho|\pi'\rangle \geq \text{OPT} - \epsilon$?

We can parametrize a product state using a vector $\vec{z} \in \mathbb{C}^n$:

$$|\pi_{\vec{z}}\rangle = \bigotimes \frac{|0\rangle + z_i|1\rangle}{\sqrt{1 + |z_i|^2}}$$

We can parametrize a product state using a vector $\vec{z} \in \mathbb{C}^n$:

$$|\pi_{\vec{z}}\rangle = \bigotimes \frac{|0\rangle + z_i|1\rangle}{\sqrt{1 + |z_i|^2}}$$

If we assume that our starting product state $|\pi\rangle = |0^n\rangle$, then being close to $|\pi\rangle$, means that $\|\vec{z}\|_2$ is small.

Observation: When $\|\vec{z}\|_2$ is small, $|\pi_{\vec{z}}\rangle$ is mostly supported on low Hamming weight strings.

Observation: When $\|\vec{z}\|_2$ is small, $|\pi_{\vec{z}}\rangle$ is mostly supported on low Hamming weight strings.

Lemma: Assume that $\|\vec{z}\|_2 \leq C$, and let $d = O(C + \log(1/\epsilon))$. Let Π_d be the projection onto all strings $|x\rangle$ with Hamming weight at most d , then

$$\|\Pi_d |\pi_{\vec{z}}\rangle\|_2 \geq 1 - \epsilon$$

Observation: When $\|\vec{z}\|_2$ is small, $|\pi_{\vec{z}}\rangle$ is mostly supported on low Hamming weight strings.

Lemma: Assume that $\|\vec{z}\|_2 \leq C$, and let $d = O(C + \log(1/\epsilon))$. Let Π_d be the projection onto all strings $|x\rangle$ with Hamming weight at most d , then

$$\|\Pi_d |\pi_{\vec{z}}\rangle\|_2 \geq 1 - \epsilon$$

In other words, if we want to (approximately) optimize $\vec{z}$, we just have to learn ρ on the small subspace Π_d .

Since $\dim(\Pi_d) = \binom{n}{d} = O(n^{C+\log(1/\epsilon)})$, we can do subspace tomography with copy and time complexity $\text{poly}(n^{C+\log(1/\epsilon)})$.

Since $\dim(\Pi_d) = \binom{n}{d} = O(n^{C+\log(1/\epsilon)})$, we can do subspace tomography with copy and time complexity $\text{poly}(n^{C+\log(1/\epsilon)})$.

The rest of our algorithm will be completely classical!

Recall: Let $|\pi\rangle$ be a product state such that $\langle\pi|\rho|\pi\rangle \geq \text{OPT} - 10\epsilon$, can we find a nearby $|\pi'\rangle$ such that $\langle\pi'|\rho|\pi'\rangle \geq \text{OPT} - \epsilon$?

Recall: Let $|\pi\rangle$ be a product state such that $\langle\pi|\rho|\pi\rangle \geq \text{OPT} - 10\epsilon$, can we find a nearby $|\pi'\rangle$ such that $\langle\pi'|\rho_d|\pi'\rangle \geq \text{OPT} - \epsilon$?

Restrict ρ to a low Hamming weight subspace



Recall: Let $|\pi\rangle$ be a product state such that $\langle\pi|\rho|\pi\rangle \geq \text{OPT} - 10\epsilon$, can we find a nearby $|\pi'\rangle$ such that $\langle\pi'|\rho_d|\pi'\rangle \geq \text{OPT} - \epsilon$?

We want to optimize over $\|\vec{z}\|_2 \leq C$:

$$\langle\pi_{\vec{z}}|\rho_d|\pi_{\vec{z}}\rangle = \frac{1}{\prod(1 + |z_i|^2)} \sum_{|x|,|x'|\leq d} \langle x|\rho_d|x'\rangle \left(\prod_{x_i=1} \bar{z}_i \prod_{x'_i=1} z_i \right)$$

Recall: Let $|\pi\rangle$ be a product state such that $\langle\pi|\rho|\pi\rangle \geq \text{OPT} - 10\epsilon$, can we find a nearby $|\pi'\rangle$ such that $\langle\pi'|\rho_d|\pi'\rangle \geq \text{OPT} - \epsilon$?

We want to optimize over $\|\vec{z}\|_2 \leq C$:

$$\langle\pi_{\vec{z}}|\rho_d|\pi_{\vec{z}}\rangle = \frac{1}{\prod(1 + |z_i|^2)} \underbrace{\sum_{|x|,|x'|\leq d} \langle x|\rho_d|x'\rangle \left(\prod_{x_i=1} \bar{z}_i \prod_{x'_i=1} z_i \right)}_{\text{A polynomial of degree } 2d \text{ in } \vec{z} \text{ ☺}}$$

Recall: Let $|\pi\rangle$ be a product state such that $\langle\pi|\rho|\pi\rangle \geq \text{OPT} - 10\epsilon$, can we find a nearby $|\pi'\rangle$ such that $\langle\pi'|\rho_d|\pi'\rangle \geq \text{OPT} - \epsilon$?

We want to optimize over $\|\vec{z}\|_2 \leq C$:

$$\langle\pi_{\vec{z}}|\rho_d|\pi_{\vec{z}}\rangle = \underbrace{\frac{1}{\prod(1 + |z_i|^2)}}_{\text{This term is annoying } \text{☹}} \sum_{|x|, |x'| \leq d} \langle x|\rho_d|x'\rangle \left(\prod_{x_i=1} \bar{z}_i \prod_{x'_i=1} z_i \right)$$

This term is annoying ☹

How do we deal with this annoying term in front?

$$\frac{1}{\prod(1 + |z_i|^2)}$$

Idea 1: When all z_i are small (order ϵ), $\frac{1}{(1+|z_i|^2)} \approx e^{-|z_i|^2}$, so if we knew they were all small, we would have:

$$\frac{1}{\prod(1 + |z_i|^2)} \approx e^{-\|z\|_2^2}$$

How do we deal with this annoying term in front?

$$\frac{1}{\prod(1 + |z_i|^2)}$$

Idea 2: Not too many z_i can be large if $\|z_i\|_2 \leq C$, just guess which z_i are large and guess their values (i.e. just brute force over them), and optimize over the remaining terms (with smaller norm).

How do we deal with this annoying term in front?

After guessing, up to re-scaling, we can say that maximizing over close product states is equivalent to finding:

$$\max_{\substack{\|\vec{z}\|_2 \leq 1 \\ \|\vec{z}\|_\infty \leq \mu}} p(\vec{z})$$

Where p is a degree $2d$ polynomial.

But... solving a polynomial optimization is not easy in general.

But... solving a polynomial optimization is not easy in general.

Luckily, our polynomial only takes on small values (since it corresponds to a fidelity between two quantum states). So we can prove the following:

Lemma (informal): There exists a subspace V of dimension $O(d^2/\epsilon)$ such that

$$|p(\vec{z}) - p(V \vec{z})| \leq \epsilon.$$

Lemma (informal): There exists a subspace V of dimension $O(d^2/\epsilon)$ such that

$$|p(\vec{z}) - p(V \vec{z})| \leq \epsilon.$$

Morally: To optimize the polynomial *without ℓ_∞ constraints*, we just have to brute force over V , taking time $O((1/\epsilon)^{\text{poly}(d,\epsilon)})$.

Lemma (informal): There exists a subspace V of dimension $O(d^2/\epsilon)$ such that

$$|p(\vec{z}) - p(V \vec{z})| \leq \epsilon.$$

Morally: To optimize the polynomial *without ℓ_∞ constraints*, we just have to brute force over V , taking time $O((1/\epsilon)^{\text{poly}(d,\epsilon)})$.

Again, we can just guess the coordinates that saturate the constraints and remove them!

Putting everything together

Full algorithm (for growing a single candidate):

1. Use subspace tomography to get a description of ρ_d .
2. Guess the coordinates of z_i that are large and their values.
3. Guess which coordinates saturate ℓ_∞ constraints exactly.
4. For the remaining coordinates, solve the ℓ_2 constrained polynomial optimization problem.

Putting everything together

Full algorithm (for growing a single candidate):

1. Use subspace tomography to get a description of ρ_d .
2. Guess the coordinates of z_i that are large and their values.
3. Guess which coordinates saturate ℓ_∞ constraints exactly.
4. For the remaining coordinates, solve the ℓ_2 constrained polynomial optimization problem.

In reality, we try to maintain a set of “all” good candidate states, and grow them all!

Lower bound

Lower bound

The lower bound comes from basically inverting the construction for the $d = 2$ subspace. Given a 4-tensor (think: degree 4 polynomial),

Lower bound

The lower bound comes from basically inverting the construction for the $d = 2$ subspace. Given a 4-tensor (think: degree 4 polynomial),

$$T \mapsto \sum_{i,j,k,\ell} T_{i,j,k,\ell} |e_i, e_j, e_k, e_\ell\rangle$$

Lower bound

The lower bound comes from basically inverting the construction for the $d = 2$ subspace. Given a 4-tensor (think: degree 4 polynomial),

$$T \mapsto \sum_{i,j,k,\ell} T_{i,j,k,\ell} |e_i, e_j, e_k, e_\ell\rangle$$

The best product state approximation, $\vec{z}$, to this* yields a solution to the tensor optimization $T(\vec{z}, \vec{z}, \vec{z}, \vec{z})$, which is NP-hard.

Open questions

- Agnostic tomography for other states?
 - Free-Fermionic states?
 - Low-degree circuits?
 - QAC_0 ?
- Tolerant testing for these models?
- Agnostic learning for other models?
 - Unitaries/Channels?
 - Hamiltonians evolution?

Thanks for listening!