Average-case quantum complexity from glassiness

Eric R. Anschuetz

eans@caltech.edu

December 9, 2025

arXiv:2505.00087 [quant-ph] and arXiv:2510.08497 [quant-ph]

What makes Gibbs state preparation hard?

What makes optimization hard?

What makes optimization hard?

(or, why complexity theory can be misleading)

Statistical-Computational Gaps

Canonical example of an NP-hard optimization problem: MAX-CLIQUE

- Given a graph on n vertices, what is its largest clique?

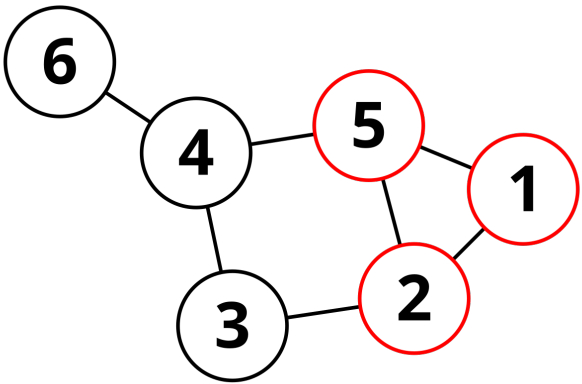

Statistical-Computational Gaps

PCP theorem: MAX-CLIQUE is even hard to *approximate*

Theorem (D. Zuckerman, Theory Comput. **3**, 103 (2007))

For any $\epsilon > 0$ *it is* NP*-hard to approximate* MAX-CLIQUE *to an approximation ratio* $\frac{1}{n^{1-\epsilon}}$

- Even a trivial clique of a single vertex achieves $\frac{1}{n}$—you can't do much better!

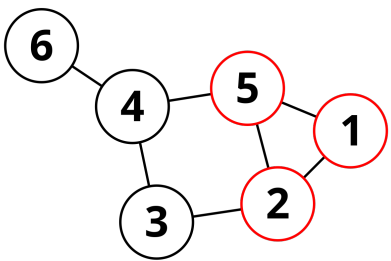

Statistical-Computational Gaps

- Complexity theory gives us hardness for *adversarially chosen* examples
- For an Erdős–Rényi graph $\sim G\left(n, \frac{1}{2}\right)$, $\mathrm{MAX\text{-}CLIQUE} \sim 2\log_2(n)$ w.h.p.
  - Trivial cliques achieve an approximation ratio $\sim \frac{1}{2\log_2(n)}$, already beating what PCP tells us for worst-case instances!

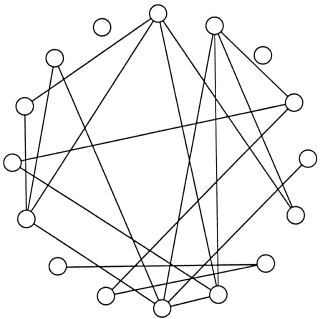

Statistical-Computational Gaps

In fact, we can do better than a factor-$2\log_2(n)$ approximation:[1]

1. Choose a random vertex v.
2. Throw away the vertices not adjacent to v.
3. Repeat.

[1] R. M. Karp, tech. rep. UCB/ERL M581 (1976)

Statistical-Computational Gaps

In fact, we can do better than a factor-$2\log_2(n)$ approximation:[1]

1. Choose a random vertex v.
2. Throw away the vertices not adjacent to v.
3. Repeat.

For a graph $\sim G\left(n, \frac{1}{2}\right)$, you throw away roughly half of the vertices each time, giving a clique of size $\sim \log_2(n)\ldots$

[1] R. M. Karp, tech. rep. UCB/ERL M581 (1976)

Statistical-Computational Gaps

In fact, we can do better than a factor-$2\log_2(n)$ approximation:[1]

1. Choose a random vertex v.
2. Throw away the vertices not adjacent to v.
3. Repeat.

For a graph $\sim G\left(n, \frac{1}{2}\right)$, you throw away roughly half of the vertices each time, giving a clique of size $\sim \log_2(n)$...

...and typically achieving an approximation ratio of $\frac{1}{2}$!

[1] R. M. Karp, tech. rep. UCB/ERL M581 (1976)

Statistical-Computational Gaps

In fact, we can do better than a factor-$2\log_2(n)$ approximation:[1]

1. Choose a random vertex v.
2. Throw away the vertices not adjacent to v.
3. Repeat.

For a graph $\sim G\left(n, \frac{1}{2}\right)$, you throw away roughly half of the vertices each time, giving a clique of size $\sim \log_2(n)\ldots$

$\ldots$ and typically achieving an approximation ratio of $\frac{1}{2}$!

A similar story holds if one is interested in sampling from the Boltzmann distribution of configurations

[1] R. M. Karp, tech. rep. UCB/ERL M581 (1976)

Statistical-Computational Gaps

Can we do *even better*?

Statistical-Computational Gaps

Can we do *even better*?

The Largest Clique of a Random Graph: The Most "Embarrassing" Open Problem in Random Structures

Imagine a club with N members, in which about 50% of the $N(N-1)/2$ member pairs know each other personally, and the remaining 50% of the members do not. You want to find a largest clique in this club, namely, the largest group of members out of the N members who all know each other. What is the typical size $c^\star$ of such a clique? How easy is it to find one? This question can

Even after 50 years, no. :(

Statistical-Computational Gaps

MAX-CLIQUE is the classic example of a *statistical-computational gap* in a maximization problem:

- *Statistically*, we know what the maximum is without doing any work!
- *Computationally*, we know of no efficient algorithms outputting the arg max

Statistical-Computational Gaps

Without complexity theory, how can we characterize problems exhibiting statistical-computational gaps?

Statistical-Computational Gaps

Without complexity theory, how can we characterize problems exhibiting statistical-computational gaps?

The Overlap Gap Property

Generally, consider an energy function $E \sim \mathcal{F}$:

$$\boldsymbol{z}^* := \underset{\boldsymbol{z} \in \{0,1\}^{\times n}}{\arg\max}\, E(\boldsymbol{z}),$$
$$E^* = E(\boldsymbol{z}^*).$$

The Overlap Gap Property

Generally, consider a combinatorial optimization problem $E \sim \mathcal{F}$:

$$\boldsymbol{z}^* := \underset{\boldsymbol{z} \in \{0,1\}^{\times n}}{\arg\max}\, E(\boldsymbol{z}),$$
$$E^* = E(\boldsymbol{z}^*).$$

Consider space of high-value points:

$$S_\mu^E := \left\{ \boldsymbol{z} \in \{0,1\}^{\times n} : E(\boldsymbol{z}) \geq \mu E^* \right\}$$

The Overlap Gap Property

E satisfies the *overlap gap property* (OGP) with parameters (μ, ν_1, ν_2) if w.h.p.:

$$\left\{(\boldsymbol{x}, \boldsymbol{y}) \in S_\mu^E \times S_\mu^E : d_{\mathsf{H}}(\boldsymbol{x}, \boldsymbol{y}) \in [\nu_1 n, \nu_2 n]\right\} = \varnothing$$

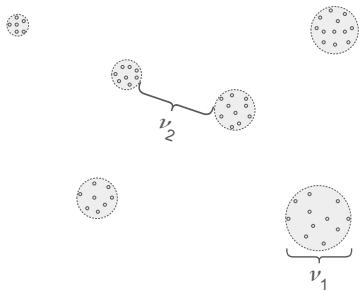

The (Ensemble) Overlap Gap Property

E satisfies the *ensemble overlap gap property* (e-OGP) with parameters (μ, ν_1, ν_2) if w.h.p.:

$$\left\{(\boldsymbol{x}, \boldsymbol{y}) \in S_\mu^E \times S_\mu^{E'} : d_{\mathsf{H}}(\boldsymbol{x}, \boldsymbol{y}) \in [\nu_1 n, \nu_2 n]\right\} = \varnothing$$

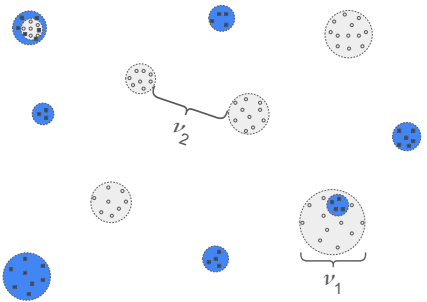

The (Ensemble) Overlap Gap Property

Why does the e-OGP obstruct algorithms?

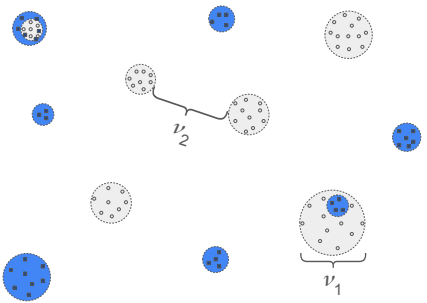

The (Ensemble) Overlap Gap Property

Why does the e-OGP obstruct algorithms?

- Independent instances have distant high-energy configurations

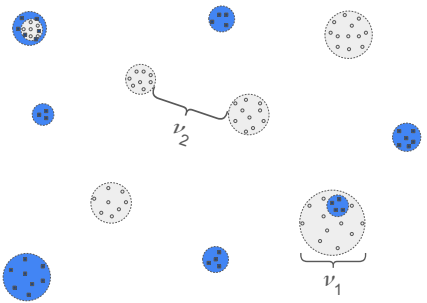

The (Ensemble) Overlap Gap Property

Why does the e-OGP obstruct algorithms?

- Independent instances have distant high-energy configurations
- By e-OGP, interpolations between independent instances have clustered solutions

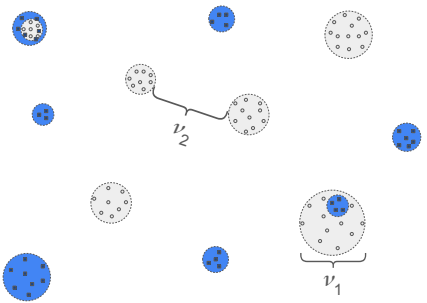

The (Ensemble) Overlap Gap Property

Why does the e-OGP obstruct algorithms?

- Independent instances have distant high-energy configurations
- By e-OGP, interpolations between independent instances have clustered solutions
- Algorithms have to "jump the gap" at some point along any interpolation path; Lipschitz algorithms can't do this!

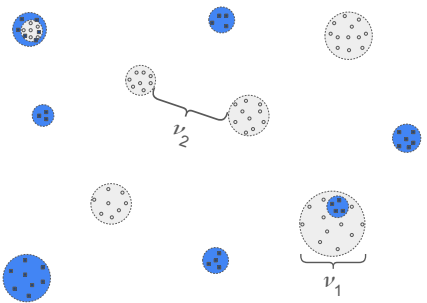

Sampling

The OGP also implies a *dynamical spin glass transition* at low temperatures[2]

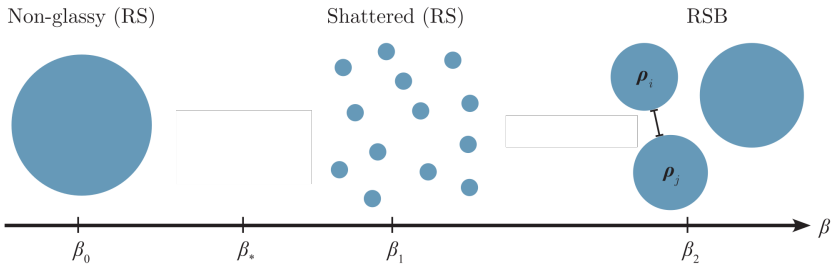

Support of Boltzmann distribution "shatters" into exponentially many states...

$\Longrightarrow$ slow mixing from a worst-case starting point due to extensive free energy wells

[2] D. Gamarnik et al., Probab. Theory Relat. Fields **193**, 89 (2025)

Classical Problems

In practice, OGP *matches* performance of best computationally efficient algorithms

Problem description	OGP matches known algorithms	References
Cliques in Erdös-Rényi graphs	Yes	Based on references below
Independent Sets in sparse Erdös-Rényi graphs	Yes	[GS17a],[RV17], [GJW20a],[Wei20],[FGG20a]
Random K-SAT	Yes	[GS17b],[COHH17],[BH21]
Largest submatrix problem	Not known	[GL18]
Matching in random hypergraphs	Not known	[CGPR19]
Ground states of spin glasses	Yes	[GJ19],[GJW20b][Sub18],[Mon19], [AMS20],
Number partitioning	Yes (up to sub-exponential factors)	[GK21]
Hidden Clique problem	Yes	[GZ19]
Sparse Linear Regression	Yes (up to a constant factor)	[GI17]
Principal submatrix recovery	Not known	[GJS19],[AWZ20]

A Quantum OGP?

Is there a *quantum* OGP for Hamiltonians?

$$\boldsymbol{H} = \sum_{E} E \left|E\right\rangle \left\langle E\right|$$

- This would tell us when we have slow mixing, easy-to-prepare low-energy states, quantum memories, ...

A Quantum OGP?

$$\boldsymbol{H} = \sum_{E} E \left|E\right\rangle \left\langle E\right|$$

Unfortunately, immediately run into issues...

A Quantum OGP?

$$\boldsymbol{H} = \sum_{E} E \left|E\right\rangle \left\langle E\right|$$

Unfortunately, immediately run into issues...

- If $|E\rangle$ and $|E'\rangle$ are high energy, so is:

$$|t\rangle := \sqrt{t}\,|E\rangle + \sqrt{1-t}\,|E'\rangle$$

for all t; how can high-energy states be "gapped"?

A Quantum OGP?

$$\boldsymbol{H} = \sum_E E \left|E\right\rangle \left\langle E\right|$$

Unfortunately, immediately run into issues...

- If $\left|E\right\rangle$ and $\left|E'\right\rangle$ are high energy, so is:

$$\left|t\right\rangle := \sqrt{t}\left|E\right\rangle + \sqrt{1-t}\left|E'\right\rangle$$

  for all t; how can high-energy states be "gapped"?
- *Of course* algorithms preparing low-energy $\left|E\right\rangle$ are inefficient; even *reading out* E can be inefficient quantumly!

How do we get circumvent this?

Efficient Learning Implies Quantum Glassiness

Eric R. Anschuetz

Average-case quantum complexity from glassiness

Alexander Zlokapa,[1,*] Bobak T. Kiani,[2,†] and Eric R. Anschuetz[3,4,‡]

The Commutation Index

- Previous work showed quantum complementarity implies the *absence* of a spin glass transition[3]...

[3] **ERA** et al., Phys. Rev. Lett. **135**, 030602 (2025)
[4] S. Chen et al., arXiv:2404.19105

The Commutation Index

- Previous work showed quantum complementarity implies the *absence* of a spin glass transition[3]...
- For:

$$\boldsymbol{H} = \frac{1}{\sqrt{m}} \sum_{k=1}^{m} J_k \boldsymbol{A}_k,$$

define the *commutation index* as:

$$\Delta_{\boldsymbol{H}} = \sup_{\||\psi\rangle\|_2=1} \frac{1}{m} \sum_{k=1}^{m} \langle\psi| \boldsymbol{A}_k |\psi\rangle^2 .$$

If $\Delta_{\boldsymbol{H}} \to 0$ asymptotically, $\boldsymbol{H}$ has no spin glass phase at constant temperature

[3] **ERA** et al., Phys. Rev. Lett. **135**, 030602 (2025)
[4] S. Chen et al., arXiv:2404.19105

The Commutation Index

- Previous work showed quantum complementarity implies the *absence* of a spin glass transition[3]...
- For:

$$\boldsymbol{H} = \frac{1}{\sqrt{m}} \sum_{k=1}^{m} J_k \boldsymbol{A}_k,$$

define the *commutation index* as:

$$\Delta_{\boldsymbol{H}} = \sup_{\||\psi\rangle\|_2 = 1} \frac{1}{m} \sum_{k=1}^{m} \langle\psi| \boldsymbol{A}_k |\psi\rangle^2 .$$

If $\Delta_{\boldsymbol{H}} \to 0$ asymptotically, $\boldsymbol{H}$ has no spin glass phase at constant temperature

- $\Delta_{\boldsymbol{H}}^{-1}$ is *exactly** the *sample complexity of learning* $\{\langle \boldsymbol{A}_k \rangle\}_k$ in a state[4]

[3] **ERA** et al., Phys. Rev. Lett. **135**, 030602 (2025)

[4] S. Chen et al., arXiv:2404.19105

The Quantum OGP

This motivates the following definition of a quantum OGP:

Definition (Quantum OGP, informal)

We say $\boldsymbol{H}$ satisfies the *quantum OGP* with parameters (μ, ν_1, ν_2) if:

1. $\Delta_{\boldsymbol{H}} \not\to 0$
2. The classical shadows estimator $\frac{1}{R}\sum_{r=1}^{R} E_{\boldsymbol{H}}(\boldsymbol{y}_r)$ exhibits a *classical* OGP with parameters (μ, ν_1, ν_2)

The Quantum OGP

This motivates the following definition of a quantum OGP:

Definition (Quantum OGP, informal)

We say $\boldsymbol{H}$ satisfies the *quantum OGP* with parameters (μ, ν_1, ν_2) if:

1. $\Delta_{\boldsymbol{H}} \not\to 0$
2. The classical shadows estimator $\frac{1}{R}\sum_{r=1}^{R} E_{\boldsymbol{H}}(\boldsymbol{y}_r)$ exhibits a *classical* OGP with parameters (μ, ν_1, ν_2)

- $E_{\boldsymbol{H}}$ is a *classical* disordered model with only a constant volume overhead iff $\Delta_{\boldsymbol{H}} \not\to 0$!

The Quantum OGP

First main result: quantum OGP satisfied by sparsified *quantum p-spin model*:

$$\boldsymbol{H}_p := \frac{1}{\sqrt{\binom{n}{p}}} \sum_{\bar{i} \in \binom{[n]}{p}} \sum_{\boldsymbol{b} \in \{1,2,3\}^p} J_{\bar{i},\boldsymbol{b}} \boldsymbol{\sigma}_{\bar{i}}^{(\boldsymbol{b})}$$

The Quantum OGP Implies Algorithmic Hardness

Second main result: the quantum OGP implies hardness for stable classical *and* quantum algorithms!

Theorem (Quantum OGP obstructs stable algorithms, informal)

Let $\mathcal{A}$ be a stable algorithm, i.e.,

$$d_{W_2}\left(\mathcal{A}(\boldsymbol{X}), \mathcal{A}(\boldsymbol{Y})\right) \leq L \left\|\boldsymbol{X} - \boldsymbol{Y}\right\|_1 \tag{1}$$

for any $L = \mathrm{O}(1)$ and $\boldsymbol{X}, \boldsymbol{Y}$ the Hamiltonian coefficients. $\mathcal{A}$ cannot achieve an approximation ratio μ if the problem class satisfies the quantum OGP with parameters $(\mu - \delta, \nu_1, \nu_2)$ for some $\delta > 0$.

Stable Algorithms

d_{W_2} is the *quantum Wasserstein distance*[5]

- Roughly, states which differ by a k-local operation differ in quantum Wasserstein distance by $O(k)$
- Generalizes Hamming distance & classical Wasserstein distance, so this generalizes classical notions of stability[6]

[5] G. De Palma et al., IEEE Trans. Inf. Theory **67**, 6627 (2021)

[6] D. Gamarnik et al., in FOCS (2022); A. El Alaoui et al., Commun. Math. Phys. **406**, 1 (2025)

Stable Algorithms

Examples of algorithms stable under quantum Wasserstein distance:

- Algorithms Lipschitz in gate complexity (quantum or classical)

Stable Algorithms

Examples of algorithms stable under quantum Wasserstein distance:

- Algorithms Lipschitz in gate complexity (quantum or classical)
- $\log(n)$-time Lindbladian evolution (from any initial state)

Stable Algorithms

Examples of algorithms stable under quantum Wasserstein distance:

- Algorithms Lipschitz in gate complexity (quantum or classical)
- $\log(n)$-time Lindbladian evolution (from any initial state)
- $\log(n)$-depth quantum neural networks (variational quantum algorithms)

Stable Algorithms

Examples of algorithms stable under quantum Wasserstein distance:

- Algorithms Lipschitz in gate complexity (quantum or classical)
- $\log(n)$-time Lindbladian evolution (from any initial state)
- $\log(n)$-depth quantum neural networks (variational quantum algorithms)
- Phase estimation to $\log(n)$ bits of precision

Stable Algorithms

Examples of algorithms stable under quantum Wasserstein distance:

- Algorithms Lipschitz in gate complexity (quantum or classical)
- $\log(n)$-time Lindbladian evolution (from any initial state)
- $\log(n)$-depth quantum neural networks (variational quantum algorithms)
- Phase estimation to $\log(n)$ bits of precision
- See paper for more!

APPENDIX B

EXAMPLES OF STABLE QUANTUM ALGORITHMS

We here relate the notion of stability in Wasserstein distance that we use in the main text to other natural notions of the stability of a quantum algorithm, as well as give explicit examples of standard quantum algorithms which are stable. As a tool to convert between various notions of stability, we will

Outlook

In conclusion, we have...

- ...constructed a way to characterize glassiness in quantum systems
- ...ruled out well-known classical *and* quantum algorithms for near-ground state preparation of disordered k-spin models
- ...connected average-case hardness to the efficiency of learning

Ideas have since already been used for other settings and algorithms

Decoded Quantum Interferometry Requires Structure

Eric R. Anschuetz,[1,2,*] David Gamarnik,[3] and Jonathan Z. Lu[4,†]

[1] *Institute for Quantum Information and Matter, Caltech, 1200 E. California Blvd., Pasadena, CA 91125, USA*
[2] *Walter Burke Institute for Theoretical Physics, Caltech, 1200 E. California Blvd., Pasadena, CA 91125, USA*
[3] *Sloan School of Management, MIT, 100 Main St., Cambridge, MA 02142, USA*
[4] *Department of Mathematics, MIT, 77 Massachusetts Ave., Cambridge, MA 02139, USA*

Average-case quantum complexity from glassiness

Alexander Zlokapa,[1,*] Bobak T. Kiani,[2,†] and Eric R. Anschuetz[3,4,‡]

[1] *Center for Theoretical Physics, MIT, 77 Massachusetts Ave., Cambridge, MA 02139, USA*
[2] *Department of Computer Science, Bowdoin College, 255 Maine Street, Brunswick, Maine 04011, USA*
[3] *Institute for Quantum Information and Matter, Caltech, 1200 E. California Blvd., Pasadena, CA 91125, USA*
[4] *Walter Burke Institute for Theoretical Physics, Caltech, 1200 E. California Blvd., Pasadena, CA 91125, USA*

Outlook

Where do we go from here?

- Can one construct general quantum "bottleneck" theorems?[7]
- Quantum algorithms achieving optimal performance? (*Algorithmic universality*[8])
- Finite dimensional or sparse models?

[7] D. Gamarnik et al., arXiv:2411.04300; T. Rakovszky et al., arXiv:2412.09598
[8] H. E. Cheairi and D. Gamarnik, arXiv:2412.18014

Thank you!

arXiv:2505.00087 [quant-ph] and arXiv:2510.08497 [quant-ph]