



Average-Case Hardness of Tensor Decomposition

Is there a range of ranks where the problem is hard classically and easy quantumly?

Omar Shehab • Three-month research proposal • Internal R&D

1. What Alex Wein has proved

Tensor decomposition is the problem of recovering the components $a_1, \dots, a_r$ from a random tensor $T = \sum_{j \leq r} a_j^{\otimes k}$. It underlies a great deal of applied machine learning—topic modeling, mixture learning, independent component analysis, phylogenetic reconstruction—and it has a long-standing gap. The components are recoverable in principle up to rank $r \sim n^{k-1}$. Every known polynomial-time algorithm stops at $r \sim n^{k/2}$. In between lies a wide band where the answer exists but nobody can compute it.

In 2023, Alexander Wein proved the first average-case hardness result for this problem [1]. Working in the low-degree-polynomial model—a restricted but powerful class of algorithms that has an excellent track record of predicting exactly where hard problems become hard—he established a sharp transition precisely at $r \sim n^{k/2}$. Below that rank, a low-degree polynomial recovers the planted component. Above it, every such polynomial fails completely, and does so even at degrees corresponding to exponential running time.

This matters because before Wein there was *no* average-case hardness evidence for tensor decomposition of any kind. It had been an open problem since 2015. His result is strong evidence that the classical algorithms we have are the best classical algorithms there are.

2. The question

Wein's theorem constrains *classical* algorithms. It says nothing whatever about quantum ones.

That leaves the question this project exists to answer:

Is there an interval of ranks in which random tensor decomposition is hard for a classical computer and easy for a quantum computer—on average, not merely in the worst case?

The emphasis on *average* case is the point. Worst-case quantum speedups for structured algebraic problems are well known and rarely useful, because real data is not adversarial. A separation on random instances would be a genuinely different claim, and tensor decomposition is an unusually good place to look for one: the problem is natural rather than contrived, the classical hardness is now on firm footing, and the instances are random by construction.

We do not know whether such an interval exists. That is the honest starting position, and it is why this is research rather than engineering.

3. The state of the art

Four papers define the current frontier, two on each side.

Classical hardness. Wein, *Average-Case Complexity of Tensor Decomposition for Low-Degree Polynomials* [1]. The result described above.

Classical algorithms. Wein, El Alaoui and Moore, *The Kikuchi Hierarchy and Tensor PCA* [2]. A hierarchy of increasingly powerful algorithms, drawn from statistical physics, that reach the best known threshold in polynomial time and trade runtime for reach beyond it. This is the classical ceiling that any quantum claim has to beat.

Quantum algorithms. Hastings, *Classical and Quantum Algorithms for Tensor Principal Component Analysis* [3], and *Accelerating Classical and Quantum Tensor PCA* [4]. The first gives a quantum algorithm for the closely related spiked-tensor problem with a fourth-power speedup over the best classical method; the second improves that to a twelfth-power speedup. These are the strongest quantum results anyone has for a problem of this type.

Taken together these four papers are the state of the art, and they leave the question open. Hastings' speedups are large but they shorten an exponential running time without making it polynomial, so on their own they do not produce the interval we are asking about. Wein's lower bound does not close the question either, because it constrains only classical algorithms. Nobody has yet built the framework that would settle it in one direction or the other.

4. What we will investigate

We will determine whether the interval exists and, if it does, where its boundaries lie. Three phases, four weeks each.

Phase 1 (weeks 1–4) — Map the classical boundary precisely. Establish where the classical algorithms of [2] actually stop for tensor decomposition, as opposed to for the spiked-tensor problem they were designed for. Wein's lower bound tells us that no efficient classical algorithm succeeds past $r \sim n^{k/2}$; we need the matching upper bound in the same setting, so that both edges of the candidate interval are pinned down rather than inferred by analogy.

Phase 2 (weeks 4–8) — Determine what the quantum algorithms actually reach. Apply the methods of [3] and [4] to tensor decomposition and establish the rank up to which they succeed, at what cost. This is where we find out whether the quantum reach exceeds the classical one anywhere, and by how much. We expect the honest answer to be a bounded improvement rather than a change of regime, and part of the work is quantifying that precisely enough to be useful.

Phase 3 (weeks 8–12) — Ask whether the interval can exist at all. The deeper question is whether Wein's argument has a quantum counterpart. If it does, it would rule the interval out and settle the matter. If it provably does not, that is itself informative: it would mean the classical hardness evidence we rely on carries no implication for quantum algorithms, and that the search is worth continuing with different tools. This phase is exploratory, and the deliverable is a well-posed formulation together with an honest account of the obstacles.

5. Method, resources, and outcomes

The work is analytical, supported by modest numerics for validation. No quantum hardware is required and none of the phases depend on hardware access. One researcher, three months, ordinary compute.

M1 (wk 4) Classical boundary established for tensor decomposition.

M2 (wk 8) Quantum reach established; the size of any advantage quantified.

M3 (wk 12) A verdict on whether the interval exists, with its location if it does; manuscript draft.

Outcomes. There are three, and all of them are worth having. If the interval exists we will have found the first average-case quantum advantage for a natural machine-learning primitive, which is a significant result. If it provably does not exist, we will have extended the reach of the strongest classical hardness framework to quantum algorithms, which is a significant result of a different kind. If neither is settled in three months, we will still have the boundaries of both the classical and quantum regions mapped for the first time, and a precise statement of what stands in the way—which is the starting point for whoever takes it further.

Risk. The main one is that Phase 3 is genuinely hard and may not converge in four weeks. It is scoped as exploratory for that reason, and Phases 1 and 2 are self-contained and deliver regardless.

References

- [1] A. S. Wein, "Average-case complexity of tensor decomposition for low-degree polynomials," arXiv:2211.05274v2 (2023).
- [2] A. S. Wein, A. El Alaoui, C. Moore, "The Kikuchi hierarchy and tensor PCA," *2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS)*, 1446–1468 (2019). arXiv:1904.03858.
- [3] M. B. Hastings, "Classical and quantum algorithms for tensor principal component analysis," *Quantum* **4**, 237 (2020). arXiv:1907.12724.

-
- [4] M. B. Hastings, “Accelerating classical and quantum tensor PCA,” arXiv:2602.10366 (2026).
 - [5] C. Bocci, L. Chiantini, G. Ottaviani, “Refined methods for the identifiability of tensors,” *Annali di Matematica Pura ed Applicata* **193**(6), 1691–1702 (2014).
 - [6] T. Ma, J. Shi, D. Steurer, “Polynomial-time tensor decompositions with sum-of-squares,” *FOCS* (2016).
 - [7] J. Ding, T. d’Orsi, C.-H. Liu, D. Steurer, S. Tiegel, “Fast algorithm for overcomplete order-3 tensor decomposition,” *Conference on Learning Theory (COLT)* (2022).
 - [8] E. Richard, A. Montanari, “A statistical model for tensor PCA,” *Advances in Neural Information Processing Systems* **27** (2014).
 - [9] P. K. Kothari, J. Xu, *Proceedings of the 2026 Annual ACM–SIAM Symposium on Discrete Algorithms (SODA)*, 2617–2632 (2026).