Large Average Subtensor Problem: Ground-State, Algorithms, and Algorithmic Barriers

Eren C. Kızıldağ (UIUC)

Algorithmic Learning Theory, 2026

Joint work with Abhishek Hegade K.R. (Princeton)

Large Average Subtensor Problem

Let $\boldsymbol{A} = (A_{i_1,\dots,i_p} : 1 \le i_1, \dots, i_p \le N) \in (\mathbb{R}^N)^{\otimes p}$ has **iid** $\mathcal{N}(0,1)$ entries and $k = \omega_N(1)$.

Algorithmic Goal:

Find (in poly-time) a $k \times \cdots \times k$ subtensor with **large average**:

$$\max_{\substack{I_1,\dots,I_p \subset [N] \\ |I_1| = \cdots = |I_p| = k}} \frac{1}{k^p} \sum_{i_1 \in I_1, \dots, i_p \in I_p} A_{i_1,\dots,i_p}$$

Focus: Ground-state, poly-time algs, algorithmic barriers.

Prior Work: Submatrix Case

- **Applications:** Genomics, biomedicine, biclustering.
  [Madeira-Oliveira 04, Shabalin-Weigman-Perou-Nobel 09, Sun-Nobel 13]

- **Stat-Comp Gap:** Optimal value is $\alpha_{\mathrm{OPT}}\sqrt{2\log N/k}$, best poly-time alg reaches $\alpha_{\mathrm{ALG}}\sqrt{2\log N/k}$, $\alpha_{\mathrm{ALG}} = \frac{4}{3} < \alpha_{\mathrm{OPT}} = \sqrt{2}$.
  Overlap Gap Property at $\alpha_{\mathrm{OGP}}\sqrt{2\log n/k}$, $\alpha_{\mathrm{OGP}} \approx 1.36$
  [Bhamidi-Dey-Nobel 13, Gamarnik-Li 16]

- Tensors ($p \geq 3$): No rigorous results for tensors.
- Matrices ($p = 2$): Rigorous results under $k = O(\log N)$.
  [Bhamidi-Dey-Nobel 13]: $k = \Theta(N)$ is an **open question**

Setup

- Focus on $k = \omega_N(1)$ with $\limsup_N k/N < 1$.
- $N \to \infty$, p is large. Formally, double limit $\lim_{p\to\infty} \lim_{N\to\infty}$.

Key Quantity

$$E_{\max} := \sqrt{\frac{2p}{k^p} \log \binom{N}{k}}$$

Ground-State Value

Theorem

$\forall \epsilon > 0$, $\exists P \in \mathbb{N}$ such that the following holds. Fix any $p \geq P$. Then, as $N \to \infty$,

$$\mathbb{P}\left[\left|\frac{M^*}{E_{\max}} - 1\right| > \epsilon\right] = 1 - \binom{N}{k}^{-\Theta(1)},$$

where $$M^* := \max_{\substack{I_1,\ldots,I_p \subset [N] \\ |I_1| = \cdots = |I_p| = k}} \frac{1}{k^p} \sum_{i_1 \in I_1, \ldots, i_p \in I_p} A_{i_1,\ldots,i_p}.$$

Proof Idea

Involved **second moment method**, use **spin glass** insights: repair 2nd moment via **Borell-TIS** inequality and an argument of [Frieze 90]

Intuition & Key Ideas

- If $Z_1, \ldots, Z_N \sim \mathcal{N}(0, \sigma^2)$, then $\max_i Z_i \approx \sigma\sqrt{2\log N}$. So, maximum of $\binom{N}{k}^p$ iid $\mathcal{N}(0, \frac{1}{k^p})$ is $\sqrt{\frac{2p}{k^p}\log\binom{N}{k}} = E_{\max}$. Tensors exhibit **correlation decay** as $p \to \infty$.
- **Multipartite Boolean p-spin glass** with **sparsity**.
- Spin glasses are difficult for **fixed** p, more tractable for **large** p [Talagrand 06, Panchenko 14, Talagrand 00, El Alaoui-Montanari-Sellke 23, Gamarnik-Jagannath-**K.** 23, El Alaoui 24, **K.** 25]

For large p, we can address $k = \Theta(N)$ through similar insights.

Poly-Time Algorithm for Matrices [Gamarnik-Li 16]

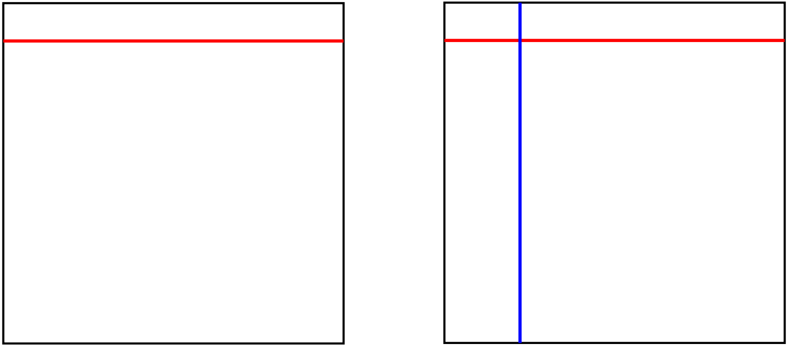

Poly-Time Algorithm for Matrices [Gamarnik-Li 16]

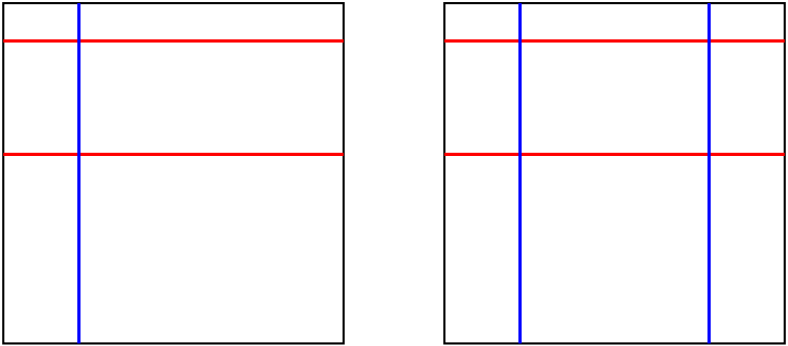

Poly-Time Algorithm for Matrices [Gamarnik-Li 16]

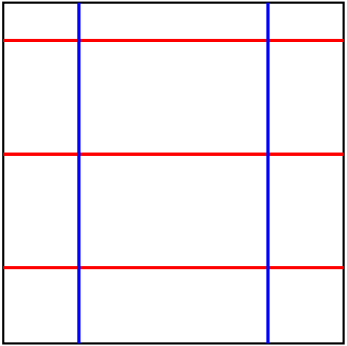

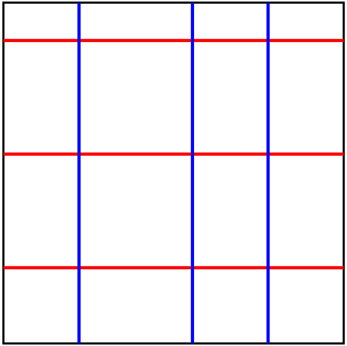

Incremental Greedy Protocol for Tensors (IGP-T)

Algorithm 1 Incremental Greedy Procedure for Tensors (IGP-T)

Input: An order-p tensor $\boldsymbol{A} \in (\mathbb{R}^N)^{\otimes p}$ and a $k \in \mathbb{N}$
Initialization (Step 1): For $1 \leq u \leq p-1$, select $i^{(u)} \in P_1$ arbitrarily, set $I_u = \{i^{(u)}\}$.
Initialization (Step 2): Let $i^{(p)} = \arg\max_{j \in P_1} A_{i^{(1)},\ldots,i^{(p-1)},j}$. Set $I_p = \{i^{(p)}\}$.
Loop-1: Repeat until $|I_1| = |I_2| = \cdots = |I_p| = k$.
Loop-2: For each $1 \leq u \leq p$, set

$$i^{(u)} = \arg\max_{j \in P_{|I_u|+1}} \sum_{\substack{i_1 \in I_1,\ldots,i_{j-1} \in I_{j-1} \\ i_{j+1} \in I_{j+1},\ldots,i_p \in I_p}} A_{i_1,\ldots,i_{j-1},j,i_{j+1},\ldots,i_p},$$

breaking ties arbitrarily. Set $I_u = I_u \cup \{i^{(u)}\}$.
Output: $I_1, \ldots, I_p$, where $I_i \subset \{1, \ldots, N\}, |I_i| = k$.

Algorithmic Guarantee

Theorem

Let IGP-T($\boldsymbol{A}$) $= (\widehat{I}_1, \ldots, \widehat{I}_p)$. Then, whp as $N \to \infty$,

$$\frac{1}{k^p} \sum_{i_1 \in \widehat{I}_1, \ldots, i_p \in \widehat{I}_p} A_{i_1,\ldots,i_p} = \big(1 + o_{k,N}(1)\big) \frac{2\sqrt{p}}{p+1} E_{\max}$$

as long as $k = o(\log^{1.5} N)$.

For $k = o(\log^{1.5} N)$, IGP-T reaches $\Theta_p(\frac{1}{\sqrt{p}})E_{\max}$ in poly-time.

[Bhamidi-Gamarnik-Gong 25]: For $k = o(N)$, IGP-T($\boldsymbol{A}$) $= O_p(\frac{1}{\sqrt{p}})E_{\max}$.

Computational Lower Bounds

Statistical-Computational Gap:

Optimal value is $E_{\max}$, algorithms reach only $E_{\max} \cdot \frac{2\sqrt{p}}{p+1}$.

Theorem

$\forall \gamma > 0$ and p sufficiently large, the model exhibits **multi OGP** *at $\gamma E_{\max}$. Consequently, stable algorithms fail to reach $\gamma E_{\max}$.*

- **Multi OGP:** Geometric property, rigorous barrier for stable algs.
  Stable Algs: Low-degree polynomials, low-depth circuits, AMP...

- As $p \to \infty$, **alg threshold** converges to zero. Similar to **REM** [Addario-Berry & Maillard 20].

Algorithmic Threshold & Future Directions

m-OGP threshold is $\gamma_p E_{\max}$ for $\gamma_p \to 0$; alg threshold is $\Theta_p(\frac{1}{\sqrt{p}})E_{\max}$.

Follow-up Work [Bhamidi-Gamarnik-Gong 25]

The onset of **Branching OGP** is precisely at $\frac{2\sqrt{p}}{p+1}E_{\max}$, for all p.

Branching OGP: Sharpest bounds for spin glasses [Huang-Sellke 21,22]

Future Work: Fixed p, algorithms for $k = \Theta(N)$.

Thank you!

OGP in Large Average Subtensor Problem

Any solutions consists of $(I_1, \ldots, I_p)$. Consider m-tuples of solutions $\mathcal{J}_i = (I_1^{(i)}, \ldots, I_p^{(i)})$

Overlap Gap Property

Largest subtensor exhibits **multi OGP** at level E if $\exists m \in \mathbb{N}$ and $\exists \beta \in (0, 1)$ such that the following holds. W.h.p. over $\boldsymbol{A}$, there is no m-tuple $\mathcal{J}_1, \ldots, \mathcal{J}_m$ satisfying:

- Every solution $\mathcal{J}_i$ is E-optimal:

$$\frac{1}{k^p} \sum_{i_1 \in I_1^{(i)}, \ldots, i_p \in I_p^{(i)}} A_{i_1, \ldots, i_p} \geq E.$$

- For any $1 \leq i < i' \leq m$ and any $1 \leq j \leq p$,

$$\frac{1}{k} \left| I_j^{(i)} \cap I_j^{(i')} \right| \approx \beta$$

Ground-State: Proof Sketch for Lower Bound

- Let N be the number of p-tuples $\mathcal{I} := (I_1, \ldots, I_p)$ with objective at least $(1-\epsilon)E_{\max}$. **Paley-Zygmund:**

$$\mathbb{P}[N \geq 1] \geq \mathbb{E}[N]^2 / \mathbb{E}[N^2].$$

- Second moment **fails**: $\mathbb{P}[N \geq 1] \geq \exp(-c^p \bar{E}^2)$, for $c < 1$ and $\bar{E} = k^{p/2}(1-\epsilon)E_{\max}$.

- M^* is **sup** of a Gaussian process. **Borell-TIS inequality**:

$$\mathbb{P}\left[\left|M^* - \mathbb{E}[M^*]\right| \geq u\right] \leq 2\exp\left(-\frac{u^2 k^p}{2}\right)$$

Repair second moment via concentration [Frieze 90].

Ground-State: Proof Sketch for Lower Bound

For all p, N large:

$$\mathbb{P}\Big[M^* \geq (1-\epsilon)E_{\max}\Big] = \mathbb{P}[N \geq 1] \\ \geq \exp\Big(-c^p k^p (1-\epsilon)^2 E_{\max}^2\Big) \\ \geq \exp\left(-\frac{k^p \epsilon^2 E_{\max}^2}{2}\right) \geq \mathbb{P}\Big[M^* \geq \mathbb{E}[M^*] + \epsilon E_{\max}\Big],$$

So, $\mathbb{E}[M^*] \geq (1-2\epsilon)E_{\max}$. Applying Borell-TIS once more gives $M^* \geq (1-3\epsilon)E_{\max}$ whp.