Asymptotically sharp bound for the column subset selection problem

Zili Xu

Math Department, Hong Kong University of Science and Technology

Joint work with Jian-feng Cai and Zhiqiang Xu

The 6th Young Scholar Symposium, May 2023
East Asia Section of Inverse Problems International Association

Outline

Outline

1. Problem setup

2. Our contributions
   - Our first result
   - Our second result

Problem setup

The column subset selection problem

The column subset selection problem (CSSP) refers the task of using the column submatrix to approximate the column space of a given matrix $\mathbf{A}$.

The column subset selection problem (CSSP)

Given a matrix $\mathbf{A} = [\mathbf{a}_1, \ldots, \mathbf{a}_d] \in \mathbb{R}^{n \times d}$ and a positive integer k, the column subset selection problem (CSSP) aims to find a subset $S \subset \{1, \ldots, d\}$ of size k, such that the approximation error $\|\mathbf{A} - \mathbf{A}_S \mathbf{A}_S^{\dagger} \mathbf{A}\|_{\xi}$ is minimized, where $\xi = 2$ or F denotes the spectral or Frobenius norm, respectively.

The column subset selection problem

The column subset selection problem (CSSP) refers the task of using the column submatrix to approximate the column space of a given matrix $\mathbf{A}$.

The column subset selection problem (CSSP)

Given a matrix $\mathbf{A} = [\mathbf{a}_1, \ldots, \mathbf{a}_d] \in \mathbb{R}^{n \times d}$ and a positive integer k, the column subset selection problem (CSSP) aims to find a subset $S \subset \{1, \ldots, d\}$ of size k, such that the approximation error $\|\mathbf{A} - \mathbf{A}_S \mathbf{A}_S^{\dagger} \mathbf{A}\|_\xi$ is minimized, where $\xi = 2$ or F denotes the spectral or Frobenius norm, respectively.

- $\mathbf{A}_S \in \mathbb{R}^{n \times k}$: the column submatrix of $\mathbf{A}$ consisting of columns indexed in the k-subset $S \subset \{1, \ldots, d\}$.
- $\mathbf{A}_S^{\dagger} \in \mathbb{R}^{k \times n}$: the Moore-Penrose pseudoinverse of $\mathbf{A}_S$
- $\mathbf{A}_S \mathbf{A}_S^{\dagger} \mathbf{A} \in \mathbb{R}^{n \times d}$: a low rank approximation to $\mathbf{A}$ by projecting all the columns of $\mathbf{A}$ to the column space of $\mathbf{A}_S$.
- $\mathbf{A} - \mathbf{A}_S \mathbf{A}_S^{\dagger} \mathbf{A} \in \mathbb{R}^{n \times d}$: the residual error matrix.

Why we consider CSSP?

Advantages of CSSP

- take advantage of the sparsity of the input matrix
- make the computed results easy to interpret in terms of the input matrix.

Applications

- machine learning
- scientific computing
- signal processing
- summarizing population genetics
- testing electronic circuits
- recommendation systems
- ...

In this talk I will mainly focus on the spectral norm version of the CSSP.

The spectral norm version of CSSP

Given a matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$ and a positive integer $k \leq \mathrm{rank}(\mathbf{A})$, we aim to find a k-subset $S \subset [d] := \{1, \ldots, d\}$ such that the approximation error $\|\mathbf{A} - \mathbf{A}_S \mathbf{A}_S^{\dagger} \mathbf{A}\|_2$ is minimized over all possible $\binom{d}{k}$ choices for the k-subsets S.

The CSSP is shown to be NP-hard.

We mainly focus on finding a k-subset S such that the approximation error $\|\mathbf{A} - \mathbf{A}_S \mathbf{A}_S^{\dagger} \mathbf{A}\|_2$ is well upper bounded.

Historical background

Historical background

Prior work mainly focuses on finding a k-subset $S \subset [d]$ such that the approximation error $\|\mathbf{A} - \mathbf{A}_S \mathbf{A}_S^{\dagger} \mathbf{A}\|_2^2$ satisfies the multiplicative bound

$$\|\mathbf{A} - \mathbf{A}_S \mathbf{A}_S^{\dagger} \mathbf{A}\|_2^2 \leq p(k, d) \cdot \|\mathbf{A} - \mathbf{A}_k\|_2^2,$$

where $p(k, d) > 1$ is a function on k and d, and $\mathbf{A}_k = \sum_{i=1}^{k} \sigma_i \mathbf{u}_i \mathbf{v}_i^{\mathrm{T}}$ is the best rank k approximation.

Historical background

- The algorithm (Gu and Eisenstat, 1996) based on rank-revealing QR decomposition gives an efficient deterministic algorithm with the multiplicative bound
$$\|\mathbf{A}-\mathbf{A}_S\mathbf{A}_S^{\dagger}\mathbf{A}\|_2^2 \leq (1+c^2k(d-k))\cdot\|\mathbf{A}-\mathbf{A}_k\|_2^2.$$
- The algorithm (Deshpande and Rademacher, 2010) based on the volume sampling, i.e., picking a subset $S \subset [d]$ with probability proportional to $\det[\mathbf{A}_S^{\mathrm{T}}\mathbf{A}_S]$, outputs a k-subset $S \subset [d]$ such that
$$\|\mathbf{A}-\mathbf{A}_S\mathbf{A}_S^{\dagger}\mathbf{A}\|_2^2 \leq (d-k)(k+1)\cdot\|\mathbf{A}-\mathbf{A}_k\|_2^2.$$
- A two-stage algorithm (Boutsidis, Drineas and Magdon-Ismail, 2014) combining RRQR based algorithms and k-leverage score sampling outputs a k-subset $S \subset [d]$ such that
$$\|\mathbf{A}-\mathbf{A}_S\mathbf{A}_S^{\dagger}\mathbf{A}\|_2^2 \leq O(k^{\frac{3}{2}}(d-k)^{\frac{1}{2}}\log k)\|\mathbf{A}-\mathbf{A}_k\|_2^2.$$
- The algorithm (Belhadji, Bardenet and Chainais, 2020) based on the projection determinantal point process outputs a k-subset $S \subset [d]$ such that
$$\|\mathbf{A}-\mathbf{A}_S\mathbf{A}_S^{\dagger}\mathbf{A}\|_2^2 \leq (1+k(\tilde{d}-k))\cdot\|\mathbf{A}-\mathbf{A}_k\|_2^2,$$
where $\tilde{d}$ is the number of the nonzero k-leverage scores.

Overall, the existing multiplicative bounds on $\|\mathbf{A}-\mathbf{A}_S\mathbf{A}_S^{\dagger}\mathbf{A}\|_2^2$ are $O(k(d-k))\cdot\|\mathbf{A}-\mathbf{A}_k\|_2^2$ and $O(k^{\frac{3}{2}}(d-k)^{\frac{1}{2}})\cdot\|\mathbf{A}-\mathbf{A}_k\|_2^2$.

The existing multiplicative bounds have the following drawbacks:

- Note that the approximation error does not exceed $\|\mathbf{A}\|_2^2$,i.e.,

$$\|\mathbf{A}-\mathbf{A}_S\mathbf{A}_S^{\dagger}\mathbf{A}\|_2^2=\|(\mathbf{I}_n-\mathbf{A}_S\mathbf{A}_S^{\dagger})\mathbf{A}\|_2^2\leq\|\mathbf{A}\|_2^2.$$

  However, the existing multiplicative bounds might be larger than $\|\mathbf{A}\|_2^2$ if $\|\mathbf{A}-\mathbf{A}_k\|_2^2=\sigma_{k+1}^2$ is large enough.
- The existing multiplicative bound are far from optimal. The lower bound on the approximation error is shown to be $\frac{d}{k}\cdot\|\mathbf{A}-\mathbf{A}_k\|_2^2$, so there is a large gap between the lower bound and the existing multiplicative bounds.

There is also a bulk of papers focus on deriving the following relative-error bounds:

$$\|\mathbf{A} - \mathbf{A}_S \mathbf{A}_S^{\dagger} \mathbf{A}\|_2 \le (1+\varepsilon) \cdot \|\mathbf{A} - \mathbf{A}_k\|_2,$$
$$\|\mathbf{A} - \mathbf{A}_S \mathbf{A}_S^{\dagger} \mathbf{A}\|_2 \le \|\mathbf{A} - \mathbf{A}_k\|_2 + \varepsilon \|\mathbf{A}\|_2,$$

where $\varepsilon > 0$ is a given error parameter.

To achieve the relative error bound, the size of S often needs to be larger than k and it is dependent on the error parameter ε.

In this talk, I mainly focus on the case when the size of S is given as input, so the discussion of the relative-error bounds is beyond the scope of this talk.

Outline

Outline

Our first result is an asymptotically sharp bound on the approximation error $\|\mathbf{A} - \mathbf{A}_S \mathbf{A}_S^\dagger \mathbf{A}\|_2^2$.

Reformulate the CSSP

For each subset $S \subset [d] = \{1, \ldots, d\}$, we define the degree d polynomial

$$p_S(x) := \det[x \cdot \mathbf{I}_d - (\mathbf{A} - \mathbf{A}_S\mathbf{A}_S^{\dagger}\mathbf{A})^{\mathrm{T}}(\mathbf{A} - \mathbf{A}_S\mathbf{A}_S^{\dagger}\mathbf{A})].$$

A simple observation is that

$$\|\mathbf{A} - \mathbf{A}_S\mathbf{A}_S^{\dagger}\mathbf{A}\|_2^2 = \text{maxroot } p_S(x).$$

Reformulate the CSSP

For each subset $S \subset [d] = \{1, \ldots, d\}$, we define the degree d polynomial

$$p_S(x) := \det[x \cdot \mathbf{I}_d - (\mathbf{A} - \mathbf{A}_S \mathbf{A}_S^{\dagger} \mathbf{A})^{\mathrm{T}} (\mathbf{A} - \mathbf{A}_S \mathbf{A}_S^{\dagger} \mathbf{A})].$$

A simple observation is that

$$\|\mathbf{A} - \mathbf{A}_S \mathbf{A}_S^{\dagger} \mathbf{A}\|_2^2 = \mathrm{maxroot}\ p_S(x).$$

Then we can reformulate the spectral norm version of CSSP as follows.

The spectral norm version of CSSP

Given a matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$ and a positive integer $k \leq \mathrm{rank}(\mathbf{A})$, we aim to find a polynomial $p_{\hat{S}}(x)$ in the set $\{p_S(x)\}_{S \subset [d], |S| = k}$, such that the largest root of $p_{\hat{S}}(x)$ is minimized.

Reformulate the CSSP

For each subset $S \subset [d] = \{1, \ldots, d\}$, we define the degree d polynomial

$$p_S(x) := \det[x \cdot \mathbf{I}_d - (\mathbf{A} - \mathbf{A}_S \mathbf{A}_S^\dagger \mathbf{A})^{\mathrm{T}} (\mathbf{A} - \mathbf{A}_S \mathbf{A}_S^\dagger \mathbf{A})].$$

A simple observation is that

$$\|\mathbf{A} - \mathbf{A}_S \mathbf{A}_S^\dagger \mathbf{A}\|_2^2 = \mathrm{maxroot}\ p_S(x).$$

Then we can reformulate the spectral norm version of CSSP as follows.

The spectral norm version of CSSP

Given a matrix $\mathbf{A} \in \mathbb{R}^{n\times d}$ and a positive integer $k \leq \mathrm{rank}(\mathbf{A})$, we aim to find a polynomial $p_{\hat{S}}(x)$ in the set $\{p_S(x)\}_{S\subset[d],|S|=k}$, such that the largest root of $p_{\hat{S}}(x)$ is minimized.

Lemma 1 (Jian-Feng Cai, Zhiqiang Xu, Zili Xu)

Let $\mathbf{A}$ be a matrix in $\mathbb{R}^{n\times d}$. Then for each positive integer $k \leq \mathrm{rank}(\mathbf{A})$, there exists a k-subset $\hat{S} \subset [d]$ such that $\mathrm{rank}(\mathbf{A}_{\hat{S}}) = k$ and

$$\mathrm{maxroot}\ p_{\hat{S}}(x) \leq \mathrm{maxroot} \sum_{S\subset[d],|S|=k} \det[\mathbf{A}_S^T \mathbf{A}_S] \cdot p_S(x).$$

We next estimate the largest root of our expected polynomial.

Lemma 2 (Jian-Feng Cai, Zhiqiang Xu, Zili Xu)

Let $\mathbf{A} \in \mathbb{R}^{n\times d}$ be a matrix of rank $t \le \min\{d, n\}$. Let λ_i be the i-th largest eigenvalue value of $\mathbf{A}^{\mathrm{T}}\mathbf{A}$. Assume that $\lambda_t < \lambda_1$. Define $\alpha := \frac{t}{\sum_{i=1}^{t} \lambda_i^{-1}}$ and $\beta := \frac{\lambda_t^{-1} - \alpha^{-1}}{\lambda_t^{-1} - \lambda_1^{-1}} \in [0, 1]$. If $\beta \cdot t \le k < t$, we have

$$\operatorname{maxroot} \sum_{S \subset [d], |S| = k} \det[\mathbf{A}_S^T \mathbf{A}_S] \cdot p_S(x) \le \frac{1}{1 + c_{\mathbf{A}} \cdot \gamma_{\mathbf{A},k}} \cdot \|\mathbf{A}\|_2^2,$$

where $c_{\mathbf{A}} := \|\mathbf{A}\|_2^2/\alpha - 1 > 0$ and $\gamma_{\mathbf{A},k} := \left(\sqrt{\frac{k}{t}} - \sqrt{\frac{\beta}{1-\beta} \cdot (1 - \frac{k}{t})}\right)^2 \in [0, 1]$.

Our main results

Theorem A (Jian-Feng Cai, Zhiqiang Xu, Zili Xu)

Let $\mathbf{A} = [\mathbf{a}_1, \ldots, \mathbf{a}_d] \in \mathbb{R}^{n\times d}$ be a matrix of rank $t \le \min\{d, n\}$. For each $1 \le i \le t$, let λ_i be the i-th largest eigenvalue value of $\mathbf{A}^\mathrm{T}\mathbf{A}$. Assume that $\lambda_t < \lambda_1$. Define $\alpha := \frac{t}{\sum_{i=1}^t \lambda_i^{-1}}$ and $\beta := \frac{\lambda_t^{-1}-\alpha^{-1}}{\lambda_t^{-1}-\lambda_1^{-1}}$. Then for any positive integer k satisfying $\beta \cdot t \le k < t$, there exists a subset $S \subset [d]$ of size k such that $\mathrm{rank}(\mathbf{A}_S) = k$ and

$$\|\mathbf{A} - \mathbf{A}_S\mathbf{A}_S^\dagger\mathbf{A}\|_2^2 \le \frac{1}{1 + c_\mathbf{A} \cdot \gamma_{\mathbf{A},k}} \cdot \|\mathbf{A}\|_2^2, \qquad (1)$$

where $c_\mathbf{A} := \|\mathbf{A}\|_2^2/\alpha - 1 > 0$ and $\gamma_{\mathbf{A},k} := (\sqrt{\frac{k}{t}} - \sqrt{\frac{\beta}{1-\beta} \cdot (1 - \frac{k}{t})})^2 \in [0, 1]$.

- $\gamma_{\mathbf{A},k}$ is increasing in k, so $\frac{1}{1+c_\mathbf{A}\cdot\gamma_{\mathbf{A},k}}$ is decreasing in k. As k increases from $\beta \cdot t$ to t, our bound in (1) decreases from $\|\mathbf{A}\|_2^2$ to α.
- The quantity α is a sharp upper bound on the minimal approximation error $\|\mathbf{A} - \mathbf{A}_S\mathbf{A}_S^\dagger\mathbf{A}\|_2^2$ for $|S| = t - 1$.
- A deterministic polynomial-time algorithm that achieves the bound in (1) is designed.

Comparison to the multiplicative bounds

Advantages:

- Our bound is strictly less than $\|\mathbf{A}\|_2^2$, while the multiplicative bounds $O(k(d-k)) \cdot \|\mathbf{A}-\mathbf{A}_k\|_2^2$ and $O(k^{\frac{3}{2}}(d-k)^{\frac{1}{2}}) \cdot \|\mathbf{A}-\mathbf{A}_k\|_2^2$ might be larger than $\|\mathbf{A}\|_2^2$.
- Our bound is asymptotically sharp when the matrix $\mathbf{A} \in \mathbb{R}^{n\times d}$ obeys a spectral power-law decay.

Overview of the proof

Proof of Lemma 1

Recall that we define $p_S(x) := \det[x \cdot \mathbf{I}_d - (\mathbf{A} - \mathbf{A}_S\mathbf{A}_S^{\dagger}\mathbf{A})^{\mathrm{T}}(\mathbf{A} - \mathbf{A}_S\mathbf{A}_S^{\dagger}\mathbf{A})]$ for each subset $S \subset [d]$.

Lemma 1 (Jian-Feng Cai, Zhiqiang Xu, Zili Xu)

Let $\mathbf{A} \in \mathbb{R}^{n\times d}$ be a matrix of rank $t \leq \min\{d, n\}$. Then for each positive integer $k \leq t$, there exists a subset $\hat{S} \subset [d]$ of size k such that $\mathrm{rank}(\mathbf{A}_{\hat{S}}) = k$ and

$$\text{maxroot } p_{\hat{S}}(x) \leq \text{maxroot} \sum_{S\subset[d],|S|=k} \det[\mathbf{A}_S^T\mathbf{A}_S] \cdot p_S(x). \tag{2}$$

Proof of Lemma 1

A key observation is the following formula

$$k! \sum_{S \subset [d], |S|=k} \det[\mathbf{A}_S^T \mathbf{A}_S] p_S(x)$$
$$= \sum_{i_1=1}^{d} \sum_{i_2 \notin \{i_1\}} \sum_{i_3 \notin \{i_1, i_2\}} \cdots \sum_{i_k \notin \{i_1, \ldots, i_{k-1}\}} \det[\mathbf{A}_{\{i_1,\ldots,i_k\}}{}^T \mathbf{A}_{\{i_1,\ldots,i_k\}}] p_{\{i_1,\ldots,i_k\}}(x).$$

Our proof of Lemma 1 is algorithmic:
Based on the above formula, we design a greedy algorithm which iteratively add columns to $\hat{S}$, and we show that the output $\hat{S}$ of our greedy algorithm satisfies (2).

Proof of Lemma 1

Our greedy algorithm (iteratively select k columns from $\mathbf{A} \in \mathbb{R}^{n \times d}$):

Step 1 Set $\hat{S} = \emptyset$. Observe that

$$
\begin{aligned}
& k! \sum_{S \subset [d], |S| = k} \det[\mathbf{A}_S^T \mathbf{A}_S] p_S(x) \\
&= \sum_{i_1=1}^{d} \sum_{i_2 \notin \{i_1\}} \sum_{i_3 \notin \{i_1, i_2\}} \cdots \sum_{i_k \notin \{i_1, \ldots, i_{k-1}\}} \det[\mathbf{A}_{\{i_1, \ldots, i_k\}}^T \mathbf{A}_{\{i_1, \ldots, i_k\}}] p_{\{i_1, \ldots, i_k\}}(x) \\
&=: \sum_{i_1=1}^{d} f_{i_1}(x).
\end{aligned}
$$

Proof of Lemma 1

Our greedy algorithm (iteratively select k columns from $\mathbf{A} \in \mathbb{R}^{n\times d}$):

Step 1 Set $\hat{S} = \emptyset$. Observe that

$$
\begin{aligned}
& k! \sum_{S\subset[d],|S|=k} \det[\mathbf{A}_S^T \mathbf{A}_S] p_S(x) \\
&= \sum_{i_1=1}^{d} \sum_{i_2\notin\{i_1\}} \sum_{i_3\notin\{i_1,i_2\}} \cdots \sum_{i_k\notin\{i_1,\ldots,i_{k-1}\}} \det[\mathbf{A}_{\{i_1,\ldots,i_k\}}^T \mathbf{A}_{\{i_1,\ldots,i_k\}}] p_{\{i_1,\ldots,i_k\}}(x) \\
&=: \sum_{i_1=1}^{d} f_{i_1}(x).
\end{aligned}
$$

We next investigate the interlacing property of the polynomials $f_1(x), \ldots, f_d(x)$.

The method of interlacing polynomials

Definition of interlacing

Let $g(x) = a_0 \cdot \prod_{i=1}^{d-1}(x - a_i)$ and $f(x) = b_0 \cdot \prod_{i=1}^{d}(x - b_i)$ be two real-rooted polynomials. We say that $g(x)$ interlaces $f(x)$ if

$$b_1 \leq a_1 \leq b_2 \leq a_2 \leq \cdots \leq a_{d-1} \leq b_d.$$

Definition of common interlacing

We say that real-rooted degree d polynomials $f_1(x), \ldots, f_m(x)$ have a common interlacing if there exists a polynomial $g(x)$ such that $g(x)$ interlaces $f_i(x)$ for each $i \in \{1, \ldots, m\}$.

Figure 1: $g(x)$ interlaces $f(x)$

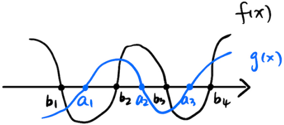

Figure 2: $f_1(x)$ and $f_2(x)$ have a common interlacing

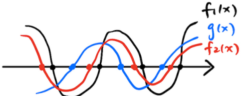

The method of interlacing polynomials

The method of interlacing polynomials is a useful tool to control the largest root of polynomials.

Theorem (Marcus-Spielman-Srivastava, 2014)

Let $\mathcal{F} = \{f_1(x), \ldots, f_m(x)\}$ be a collection of real-rooted polynomials with the same degree and positive leading coefficients.
If $f_1(x), \ldots, f_m(x)$ have a common interlacing, then there exists an integer $i \in \{1, \ldots, m\}$ such that

$$\mathrm{maxroot}\, f_i(x) \leq \mathrm{maxroot} \sum_{i=1}^{m} f_i(x). \tag{3}$$

- The above theorem basically says that the largest root of the expected polynomial provides an upper bound on the smallest largest root of $f_i(x)$.
- If $f_1(x), \ldots, f_m(x)$ fail to have a common interlacing, then (3) may not hold. For example, consider the case where $f_1(x) = (x+5)(x-9)(x-10)$ and $f_2(x) = (x+6)(x-1)(x-8)$. The largest root of $f_1(x) + f_2(x)$ is approximately 7.4.

Proof of Lemma 1

Our greedy algorithm (iteratively select k columns from $\mathbf{A} \in \mathbb{R}^{n \times d}$):

Step 1: $\hat{S} = \emptyset$. Observe that

$$k! \sum_{S \subset [d], |S|=k} \det[\mathbf{A}_S^T \mathbf{A}_S] p_S(x) = \sum_{i_1=1}^{d} \sum_{i_2 \notin \{i_1\}} \sum_{i_3 \notin \{i_1, i_2\}} \cdots \sum_{i_k \notin \{i_1, \ldots, i_{k-1}\}} \det[\mathbf{A}_{\{i_1, \ldots, i_k\}}^T \mathbf{A}_{\{i_1, \ldots, i_k\}}] p_{\{i_1, \ldots, i_k\}}(x)$$

$$=: \sum_{i_1=1}^{d} f_{i_1}(x).$$

We prove that the polynomials $f_1(x), \ldots, f_d(x)$ have a common interlacing.

Proof of Lemma 1

Our greedy algorithm (iteratively select k columns from $\mathbf{A} \in \mathbb{R}^{n \times d}$):

Step 1: $\hat{S} = \emptyset$. Observe that

$$k! \sum_{S \subset [d], |S|=k} \det[\mathbf{A}_S^T \mathbf{A}_S] p_S(x) = \sum_{i_1=1}^{d} \sum_{i_2 \notin \{i_1\}} \sum_{i_3 \notin \{i_1, i_2\}} \cdots \sum_{i_k \notin \{i_1, \ldots, i_{k-1}\}} \det[\mathbf{A}_{\{i_1, \ldots, i_k\}}^T \mathbf{A}_{\{i_1, \ldots, i_k\}}] p_{\{i_1, \ldots, i_k\}}(x)$$

$$=: \sum_{i_1=1}^{d} f_{i_1}(x).$$

We prove that the polynomials $f_1(x), \ldots, f_d(x)$ have a common interlacing.

Then there exists an integer $j_1 \in [d]$ such that

$$\text{maxroot } f_{j_1}(x) \leq \text{maxroot} \sum_{i_1=1}^{d} f_{i_1}(x) = \text{maxroot} \sum_{S \subset [d], |S|=k} \det[\mathbf{A}_S^T \mathbf{A}_S] p_S(x).$$

Then we set $\hat{S} = \{j_1\}$.

Proof of Lemma 1

Step 2: $\hat{S} = \{j_1\}$. Observe that

$$f_{j_1}(x) = \sum_{i_2 \notin \{j_1\}} \sum_{i_3 \notin \{j_1, i_2\}} \cdots \sum_{i_k \notin \{j_1, i_2, \ldots, i_{k-1}\}} \det[\mathbf{A}^T_{\{j_1, i_2, \ldots, i_k\}} \mathbf{A}_{\{j_1, i_2, \ldots, i_k\}}] p_{\{j_1, i_2, \ldots, i_k\}}(x)$$

$$=: \sum_{i_2 \notin \{j_1\}} g_{i_2}(x).$$

Proof of Lemma 1

Step 2: $\hat{S} = \{j_1\}$. Observe that

$$f_{j_1}(x) = \sum_{i_2 \notin \{j_1\}} \sum_{i_3 \notin \{j_1, i_2\}} \cdots \sum_{i_k \notin \{j_1, i_2, \ldots, i_{k-1}\}} \det[\mathbf{A}^T_{\{j_1, i_2, \ldots, i_k\}} \mathbf{A}_{\{j_1, i_2, \ldots, i_k\}}] p_{\{j_1, i_2, \ldots, i_k\}}(x)$$

$$=: \sum_{i_2 \notin \{j_1\}} g_{i_2}(x).$$

We prove that the polynomials $g_{i_2}(x), \forall\, i_2 \notin \{i_1\}$ have a common interlacing.

Proof of Lemma 1

Step 2: $\hat{S} = \{j_1\}$. Observe that

$$f_{j_1}(x) = \sum_{i_2 \notin \{j_1\}} \sum_{i_3 \notin \{j_1, i_2\}} \cdots \sum_{i_k \notin \{j_1, i_2, \ldots, i_{k-1}\}} \det[\mathbf{A}^T_{\{j_1, i_2, \ldots, i_k\}} \mathbf{A}_{\{j_1, i_2, \ldots, i_k\}}] p_{\{j_1, i_2, \ldots, i_k\}}(x)$$

$$=: \sum_{i_2 \notin \{j_1\}} g_{i_2}(x).$$

We prove that the polynomials $g_{i_2}(x), \forall\, i_2 \notin \{i_1\}$ have a common interlacing.

Then there exists $j_2 \notin \{i_1\}$ such that

$$\text{maxroot}\, g_{j_2}(x) \leq \text{maxroot} \sum_{i_2 \notin \{i_1\}} g_{i_2}(x) = \text{maxroot}\, f_{j_1}(x) \overset{\text{by Step 1}}{\leq} \text{maxroot} \sum_{S \subset [d], |S| = k} \det[\mathbf{A}_S^T \mathbf{A}_S] p_S(x).$$

Then we set $\hat{S} = \{j_1, j_2\}$.

Repeat the above argument for another $k - 2$ times.

Proof of Lemma 1

Step k: $\hat{S} = \{j_1, \ldots, j_{k-1}\}$. We prove that the polynomials

$$\det[\mathbf{A}^T_{\{j_1,\ldots,j_{k-1},i_k\}}\mathbf{A}_{\{j_1,\ldots,j_{k-1},i_k\}}]p_{\{j_1,\ldots,j_{k-1},i_k\}}(x), \forall\, i_k \notin \{j_1, \ldots, j_{k-1}\}$$

have a common interlacing.

Proof of Lemma 1

Step k: $\hat{S} = \{j_1, \ldots, j_{k-1}\}$. We prove that the polynomials

$$\det[\mathbf{A}^T_{\{j_1,\ldots,j_{k-1},i_k\}}\mathbf{A}_{\{j_1,\ldots,j_{k-1},i_k\}}]p_{\{j_1,\ldots,j_{k-1},i_k\}}(x), \forall\, i_k \notin \{j_1, \ldots, j_{k-1}\}$$

have a common interlacing.

Then there exists $j_k \notin \{j_1, \ldots, j_{k-1}\}$ such that

$$\text{maxroot } p_{\{j_1,\ldots,j_{k-1},j_k\}}(x) \leq \text{maxroot} \sum_{i_k \notin \{j_1,\ldots,j_{k-1}\}} \det[\mathbf{A}^T_{\{j_1,\ldots,j_{k-1},i_k\}}\mathbf{A}_{\{j_1,\ldots,j_{k-1},i_k\}}]p_{\{j_1,\ldots,j_{k-1},i_k\}}(x)$$

$$\overset{\text{by Step k-1}}{\leq} \cdots \overset{\text{by Step 1}}{\leq} \text{maxroot} \sum_{S\subset[d],|S|=k} \det[\mathbf{A}_S^T\mathbf{A}_S]p_S(x).$$

Then we set $\hat{S} = \{j_1, j_2, \ldots, j_k\}$. This finishes the proof of Lemma 1.

Polynomial Selection Problem

The method of interlacing polynomials is a useful tool for the Polynomial Selection Problem.

Polynomial Selection Problem

Let $\mathcal{F} = \{f_1(x), \ldots, f_m(x)\}$ be a collection of degree d polynomials with positive leading coefficients. Each polynomial in $\mathcal{F}$ is real-rooted, i.e., all of the coefficients and roots are real numbers. We aim to select a polynomial $f_i(x)$ from $\mathcal{F}$ such that the largest root of $f_i(x)$ is as small as possible.

When m is large, it is very inefficient to calculate the largest root of each polynomial and then compare them to choose the smallest one.

Polynomial Selection Problem

The method of interlacing polynomials is a useful tool for the Polynomial Selection Problem.

Polynomial Selection Problem

Let $\mathcal{F} = \{f_1(x), \ldots, f_m(x)\}$ be a collection of degree d polynomials with positive leading coefficients. Each polynomial in $\mathcal{F}$ is real-rooted, i.e., all of the coefficients and roots are real numbers. We aim to select a polynomial $f_i(x)$ from $\mathcal{F}$ such that the largest root of $f_i(x)$ is as small as possible.

When m is large, it is very inefficient to calculate the largest root of each polynomial and then compare them to choose the smallest one.

Example 1: Spectral norm version of CSSP

Let $\mathcal{F} = \{p_S(x)\}_{S\subset[d],|S|=k}$, where

$$p_S(x) := \det[x \cdot \mathbf{I}_d - (\mathbf{A} - \mathbf{A}_S\mathbf{A}_S^\dagger\mathbf{A})^{\mathrm{T}}(\mathbf{A} - \mathbf{A}_S\mathbf{A}_S^\dagger\mathbf{A})].$$

We aim to find a polynomial $p_{\hat{S}}(x)$ in $\mathcal{F}$, such that the largest root of $p_{\hat{S}}(x)$ is minimized.

Polynomial Selection Problem

Example 2: Matrix Spencer Conjecture

For all matrices $\mathbf{A}_1, \ldots, \mathbf{A}_n \in \mathbb{R}^{d \times d}$ with $\|\mathbf{A}_i\|_2 \leq 1$, there exists $\varepsilon = (\varepsilon_1, \ldots, \varepsilon_n) \in \{\pm 1\}^n$ such that $\| \sum_{i=1}^n \varepsilon_i \mathbf{A}_i \|_2 \leq O(\sqrt{n \log(d/n)})$.

An equivalent statement is as follows.

Let $\mathcal{F} = \{f_\varepsilon(x)\}_{\varepsilon \in \{\pm 1\}^n}$, where

$$f_\varepsilon(x) := \det[x \cdot \mathbf{I}_d - (\sum_{i=1}^n \varepsilon_i \mathbf{A}_i)^{\mathrm{T}} (\sum_{i=1}^n \varepsilon_i \mathbf{A}_i)].$$

There exists a polynomial $f_{\hat{\varepsilon}}(x)$ in $\mathcal{F}$, such that the largest root of $f_{\hat{\varepsilon}}(x)$ is $O(n \log(d/n))$.

Recently some special cases of Matrix Spencer Conjecture is proved by using the method of interlacing polynomials.

Proof of Lemma 2: Estimate the largest root of the expected polynomial

Lemma 3

Let $\mathbf{A}$ be a matrix in $\mathbb{R}^{n\times d}$. For any nonnegative integer k, we have

$$k! \sum_{S\subset[d],|S|=k} \det[\mathbf{A}_S^T\mathbf{A}_S]\cdot p_S(x) = (-1)^k\cdot\mathcal{R}_d\cdot\partial_x^k\cdot\mathcal{R}_d \det[x\cdot\mathbf{I}_d - \mathbf{A}^T\mathbf{A}],$$

where operator $\mathcal{R}_d$ is defined as $\mathcal{R}_d\, f(x) := x^d\cdot f(\frac{1}{x})$ for any degree d polynomial $f(x)$.

Proof of Lemma 2: Estimate the largest root of the expected polynomial

Lemma 3

Let $\mathbf{A}$ be a matrix in $\mathbb{R}^{n\times d}$. For any nonnegative integer k, we have

$$k! \sum_{S\subset[d],|S|=k} \det[\mathbf{A}_S^T\mathbf{A}_S]\cdot p_S(x) = (-1)^k\cdot\mathcal{R}_d\cdot\partial_x^k\cdot\mathcal{R}_d\ \det[x\cdot\mathbf{I}_d-\mathbf{A}^T\mathbf{A}],$$

where operator $\mathcal{R}_d$ is defined as $\mathcal{R}_d\, f(x) := x^d\cdot f(\frac{1}{x})$ for any degree d polynomial $f(x)$.

Observe that, if all roots of a degree d polynomial $f(x)$ are positive, then

$$\text{maxroot } \mathcal{R}_d\, f(x) = \frac{1}{\text{minroot } f(x)} \tag{4}$$

Combining Lemma 3 and (4), we obtain

$$\text{maxroot} \sum_{S\subset[d],|S|=k} \det[\mathbf{A}_S^T\mathbf{A}_S]\cdot p_S(x) = \frac{1}{\text{minroot } \partial_x^k\cdot\mathcal{R}_d\ \det[x\cdot\mathbf{I}_d-\mathbf{A}^T\mathbf{A}]}$$

Proof of Lemma 2

The following lemma shows how the largest root of a univariate polynomial shrink after taking derivatives.

Lemma (Ravichandran 2020)

Assume that $p(x) = \prod_{i=1}^{t}(x - \lambda_i)$ is a real-rooted polynomial of degree t, where $\lambda_i \in [0, 1]$ for each $i \in [t]$. Let $\gamma = \frac{1}{t}\sum_{i=1}^{t}\lambda_i$. Then, for each $k \geq \gamma \cdot t$,

$$\mathrm{maxroot}\ \partial_x^k\, p(x) \leq \left(\sqrt{\gamma \cdot \frac{k}{t}} + \sqrt{(1-\gamma)\cdot\left(1 - \frac{k}{t}\right)}\right)^2.$$

With some slight modification, we are able to use this lemma to estimate $\mathrm{minroot}\ \partial_x^k \cdot \mathcal{R}_d\ \det[x \cdot \mathbf{I}_d - \mathbf{A}^T\mathbf{A}]$ and prove Lemma 2.

Outline

1. Problem setup
2. Our contributions
   - Our first result
   - Our second result

Our second result

Our second results

We show that our method also works in the context of the column partition problem .

The column partition problem

Given a matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$ and a positive integer $r \geq 2$, the column partition problem aims to find an r partition $S_1 \sqcup \cdots \sqcup S_r = [d]$ such that the following residual

$$\max_{1 \leq i \leq r} \|\mathbf{A} - \mathbf{A}_{S_i^C} \mathbf{A}_{S_i^C}^{\dagger} \mathbf{A}\|_2$$

is minimized over all possible r partitions of $[d]$. Here, we use S^C to denote the complement of a subset S of $[d]$.

Our results

We show that our method also works in the context of the column partition problem .

Theorem B

Assume that $n \geq d > 1$. Let $\mathbf{A} \in \mathbb{R}^{n \times d}$ be a matrix of rank d. For each $j \in \{1, \ldots, d\}$, let λ_j be the j-th largest eigenvalue value of $\mathbf{A}^\top \mathbf{A}$. Define

$$\alpha := \max_{1 \leq j \leq d} \mathbf{a}_j^\top (\mathbf{I}_n - \mathbf{A}_{\{1,\ldots,d\}\setminus\{j\}} \mathbf{A}^\dagger_{\{1,\ldots,d\}\setminus\{j\}}) \mathbf{a}_j \quad \text{and} \quad \beta := \frac{\lambda_d^{-1} - \alpha^{-1}}{\lambda_d^{-1} - \lambda_1^{-1}}.$$

Then for any integer $r \geq 2$ satisfying that $\beta \leq \frac{(r-1)^2}{r^2}$, there exists an r partition $S_1 \sqcup \cdots \sqcup S_r = \{1, \ldots, d\}$ such that

$$\|\mathbf{A} - \mathbf{A}_{S_i^C} \mathbf{A}^\dagger_{S_i^C} \mathbf{A}\|_2^2 \leq \frac{1}{\gamma_{\mathbf{A},r} + (1 - \gamma_{\mathbf{A},r}) \cdot \frac{\lambda_1}{\lambda_d}} \cdot \|\mathbf{A}\|_2^2, \quad \forall\, i \in [r], \tag{5}$$

where $\gamma_{\mathbf{A},r} := (\sqrt{\frac{1}{r} - \frac{\beta}{r-1}} + \sqrt{\beta})^2 \in [1/r, 1]$.

Thanks for your attention!