

A Maximum likelihood estimator

Let $\mathcal{Y} \in \mathbb{R}^{n_1 \times \dots \times n_K}$ be a noisy observed tensor generated as follows:

$$\mathcal{Y} = \mathcal{X}^* + \sigma \mathcal{E} = \sum_{r=1}^R \beta_r \mathbf{u}_r^{(1)} \circ \dots \circ \mathbf{u}_r^{(K)} + \sigma \mathcal{E},$$

where \mathcal{E} is a noisy tensor whose entries are i.i.d. normal $\mathcal{N}(0, 1)$.

Let $\hat{\mathcal{X}}_{\text{MLE}}$ be the (intractable) estimator defined as

$$\hat{\mathcal{X}}_{\text{MLE}} = \arg \min_{\mathcal{X}} \left(\|\mathcal{Y} - \mathcal{X}\|_F^2 : \text{rank}(\mathcal{X}) \leq R \right).$$

We have the following performance guarantee for $\hat{\mathcal{X}}_{\text{MLE}}$:

Theorem 3. *Let $R \leq \min_k n_k/2$. Then there is a constant c such that*

$$\|\hat{\mathcal{X}}_{\text{MLE}} - \mathcal{X}^*\|_F \leq c\sigma \sqrt{R^K \sum_{k=1}^K n_k \log(2K/K_0) + \log(2/\delta)},$$

with probability at least $1 - \delta$, where $K_0 = \log(3/2)$.

Note that the factor R^K in the square root is rather conservative. In the best case, this factor reduces to linear in R and this is what we present in Section 4 as ‘‘optimistic’’ ignoring constants and δ ; see Eq. (11).

Proof of Theorem 3. Since $\hat{\mathcal{X}}_{\text{MLE}}$ is a minimizer and \mathcal{X}^* is also feasible, we have

$$\|\mathcal{Y} - \hat{\mathcal{X}}_{\text{MLE}}\|_F^2 \leq \|\mathcal{Y} - \mathcal{X}^*\|_F^2,$$

which implies

$$\begin{aligned} \|\mathcal{X}^* - \hat{\mathcal{X}}_{\text{MLE}}\|_F^2 &\leq \sigma \langle \mathcal{E}, \hat{\mathcal{X}}_{\text{MLE}} - \mathcal{X}^* \rangle \\ &\leq \sigma \|\mathcal{E}\|_{\text{op}} \|\hat{\mathcal{X}}_{\text{MLE}} - \mathcal{X}^*\|_{\text{nuc}}, \end{aligned}$$

where

$$\begin{aligned} \|\mathcal{X}\|_{\text{op}} &:= \sup_{\mathbf{u}^{(1)}, \dots, \mathbf{u}^{(K)}} \left\{ \sum_{i_1, i_2, \dots, i_K} \mathcal{X}_{i_1, i_2, \dots, i_K} u_{i_1}^{(1)} u_{i_2}^{(2)} \dots u_{i_K}^{(K)} : \right. \\ &\quad \left. \|\mathbf{u}^{(1)}\| = \|\mathbf{u}^{(2)}\| = \dots = \|\mathbf{u}^{(K)}\| = 1 \right\} \end{aligned}$$

is the tensor spectral norm and the nuclear norm

$$\begin{aligned} \|\mathcal{X}\|_{\text{nuc}} &:= \inf_{\mathbf{u}^{(1)}, \dots, \mathbf{u}^{(K)}} \left\{ \sum_r \|\mathbf{u}_r^{(1)}\| \cdot \|\mathbf{u}_r^{(2)}\| \dots \|\mathbf{u}_r^{(K)}\| : \right. \\ &\quad \left. \mathcal{X} = \sum_{r=1}^R \mathbf{u}_r^{(1)} \circ \dots \circ \mathbf{u}_r^{(K)} \right\} \end{aligned}$$

is the dual of the spectral norm.

Since both $\hat{\mathcal{X}}_{\text{MLE}}$ and \mathcal{X}^* are rank at most R , the difference $\hat{\mathcal{X}}_{\text{MLE}} - \mathcal{X}^*$ is rank at most $2R$. Moreover, any rank- R CP decomposition with $R \leq \min_k n_k$ can be reduced to an orthogonal CP decomposition with rank at most R^K via the Tucker decomposition [14]. Thus, denoting this orthogonal decomposition by $\hat{\mathcal{X}}_{\text{MLE}} - \mathcal{X}^* = \sum_{r=1}^{R^K} \tilde{\mathbf{u}}_r^{(1)} \circ \dots \circ \tilde{\mathbf{u}}_r^{(K)}$ and using $\beta_r := \|\tilde{\mathbf{u}}_r^{(1)}\| \dots \|\tilde{\mathbf{u}}_r^{(K)}\|$, we have

$$\begin{aligned} \|\hat{\mathcal{X}}_{\text{MLE}} - \mathcal{X}^*\|_{\text{nuc}} &\leq \sum_{r=1}^{R^K} \beta_r \leq \sqrt{R^K} \sqrt{\sum_{r=1}^{R^K} \beta_r^2} \\ &= \sqrt{R^K} \|\hat{\mathcal{X}}_{\text{MLE}} - \mathcal{X}^*\|_F, \end{aligned}$$

where the last equality follows because the decomposition is orthogonal.

Finally applying the tail bound for the spectral norm $\|\mathcal{E}\|_{\text{op}}$ of random Gaussian tensor \mathcal{E} [24], we obtain what we wanted. \square

Algorithm 1: Tensor denoising via the subspace norm

Input: noisy tensor \mathcal{Y} , subspace dimension H , regularization constant λ
for $k = 1$ **to** K **do**
 $\widehat{\mathbf{P}}^{(k)} \leftarrow$ top H left singular vectors of $\mathbf{Y}^{(k)}$
end for
for $k = 1$ **to** K **do**
 $\mathbf{S}^{(k)} \leftarrow \widehat{\mathbf{P}}^{(1)} \otimes \dots \otimes \widehat{\mathbf{P}}^{(k-1)} \otimes \widehat{\mathbf{P}}^{(k+1)} \otimes \dots \otimes \widehat{\mathbf{P}}^{(K)}$
end for
Output: $\widehat{\mathcal{X}} = \arg \min_{\mathcal{X}} \frac{1}{2} \|\mathcal{Y} - \mathcal{X}\|_F^2 + \lambda \|\mathcal{X}\|_s$.

B Details of optimization

For solving problem (8), we follow the alternating direction method of multipliers described in [25]. We scale the objective function in (8) by $1/\lambda$, and consider the dual problem

$$\begin{aligned} \min_{\mathcal{D}, \{\mathbf{W}^{(k)}\}_{k=1}^K} \quad & \frac{\lambda}{2} \|\mathcal{D}\|_F^2 - \langle \mathcal{D}, \mathcal{Y} \rangle \\ \text{s.t.} \quad & \max_k \|\mathbf{W}^{(k)}\| \leq 1, \\ & \mathbf{W}^{(k)} = \mathbf{D}_{(k)} \mathbf{S}^{(k)}, \quad k = 1, \dots, K, \end{aligned} \tag{12}$$

where $\mathcal{D} \in \mathbb{R}^{n_1 \times n_2 \times \dots \times n_K}$ is the dual tensor that corresponds to the residual in the primal problem (8), and $\mathbf{W}^{(k)}$'s are auxiliary variables introduced to make the problem equality constrained.

The augmented Lagrangian function of problem (12) could be written as follows:

$$\begin{aligned} & L_\eta(\mathcal{D}, \{\mathbf{W}^{(k)}\}_{k=1}^K, \{\mathbf{M}^{(k)}\}_{k=1}^K) \\ &= \frac{\lambda}{2} \|\mathcal{D}\|_F^2 - \langle \mathcal{D}, \mathcal{Y} \rangle + \sum_{k=1}^K \langle \mathbf{M}^{(k)}, \mathbf{D}_{(k)} \mathbf{S}^{(k)} - \mathbf{W}^{(k)} \rangle \\ &+ \frac{\eta}{2} \|\mathbf{D}_{(k)} \mathbf{S}^{(k)} - \mathbf{W}^{(k)}\|_F^2 + \mathbf{1}_{\|\cdot\| \leq 1}(\mathbf{W}^{(k)}), \end{aligned}$$

where $\mathbf{M}^{(k)}$'s are the multipliers, η is the augmenting parameter, and $\mathbf{1}_{\|\cdot\| \leq 1}$ is the indicator function of the unit spectral norm ball.

We follow the derivation in [25] and conclude that the updates of \mathcal{D} , $\mathbf{M}^{(k)}$ and $\mathbf{W}^{(k)}$ can be computed in closed forms. We further combine the updates of $\mathbf{W}^{(k)}$ and other steps so that it needs not to be explicitly computed. The sum of the products of $\mathbf{M}^{(k)}$ and $\mathbf{S}^{(k)\top}$ finally converges to the solution of the primal problem (8), see Algorithm 2.

The update for the Lagrangian multipliers $\mathbf{M}^{(k)}$ ($k = 1, \dots, K$) is written as singular value soft-thresholding operator defined as

$$\text{prox}_\eta^{\text{tr}}(\mathbf{Z}) = \mathbf{P} \max(\boldsymbol{\Sigma} - \eta, 0) \mathbf{Q}^\top,$$

where $\mathbf{Z} = \mathbf{P} \boldsymbol{\Sigma} \mathbf{Q}^\top$ is the SVD of \mathbf{Z} .

A notable property of the subspace norm is the computational efficiency. The update of $\mathbf{M}^{(k)}$ requires singular value decomposition, which usually dominates the costs of computation. For problem (12), the size of $\mathbf{M}^{(k)}$ is only $n_k \times H^{K-1}$. Comparing with previous approaches, e.g. the latent approach whose multipliers are $n_k \times \prod_{k' \neq k} n_{k'}$ matrices, the size of our variables is much smaller, so the per-iteration cost is reduced.

C Additional experiments

We report the experimental results when the input rank of CP and the subspace approach is are over-specified, on the same synthetic dataset as Section 4. We consider the case where the input rank is 8.

Algorithm 2: ADMM for subspace norm minimization

Input: \mathcal{Y} , λ , $\mathbf{S}^{(1)}, \dots, \mathbf{S}^{(K)}$, η , initializations $\mathcal{D}_0, \{M_0^{(1)}, \dots, M_0^{(K)}\}$
 $t = 0$

repeat

$$\mathcal{D}_{t+1} = \frac{1}{\lambda + \eta K} \left(\mathcal{Y} + K\eta\mathcal{D}_t - \sum_k \text{fold}_k \left((2M_t^{(k)} - M_{t-1}^{(k)})\mathbf{S}^{(k)\top} \right) \right)$$

for $k = 1$ **to** K **do**

$$M_{t+1}^{(k)} = \text{prox}_{\eta}^{tr} \left(M_t^{(k)} + \eta\mathcal{D}_{(k),t+1}\mathbf{S}^{(k)} \right)$$

end for

$t \leftarrow t + 1$

until convergence

Output: $\hat{\mathcal{X}} = \sum_{k=1}^K M_t^{(k)}\mathbf{S}^{(k)\top}$.

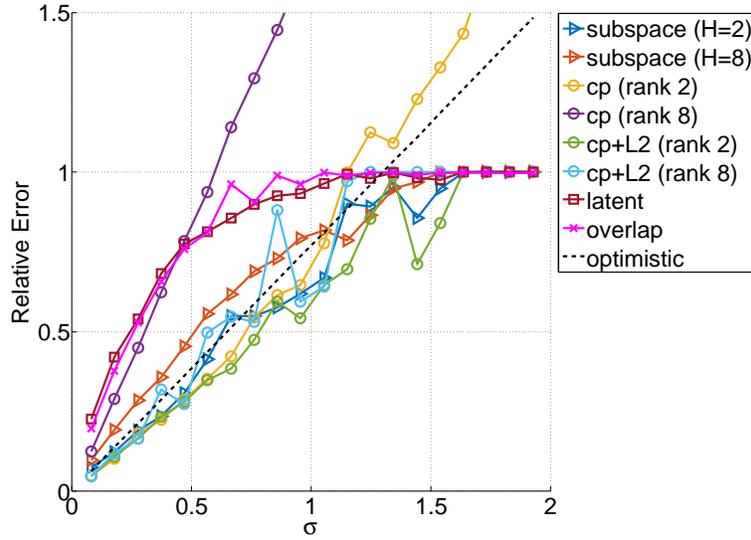


Figure 3: Tensor denoising on synthetic dataset when the input rank is larger than the truth.

We impose the ℓ_2 regularizations on the factors of CP. We test 20 values that are logarithmically spaced between 0.01 and 10 as the regularization parameter. For each value, we compute 20 solutions with random initializations and select the one with lowest objective value.

For the subspace approach, we computed solutions for 20 values of the regularization parameter that are logarithmically spaced between 1 and 1000.

As before, we report the minimum relative error obtained by the same method. The results are shown in Figure 3. We include the case the rank is specified incorrectly for comparison. Clearly, even if the rank is much larger than the truth, the subspace approach and CP are robust with proper regularization.

D Proofs

D.1 Proof of Theorem 1

We consider the second moment of $\tilde{\mathbf{X}}$:

$$\begin{aligned}\tilde{\mathbf{X}}\tilde{\mathbf{X}}^\top &= \beta^2\mathbf{u}\mathbf{u}^\top + \sigma^2\mathbf{E}\mathbf{E}^\top + \beta\sigma(\mathbf{u}\mathbf{v}^\top\mathbf{E}^\top + \mathbf{E}\mathbf{v}\mathbf{u}^\top) \\ &= \underbrace{\beta^2\mathbf{u}\mathbf{u}^\top + m\sigma^2\mathbf{I}}_{\mathbf{B}} + \\ &\quad \underbrace{\sigma^2\mathbf{E}\mathbf{E}^\top - m\sigma^2\mathbf{I} + \beta\sigma(\mathbf{u}\mathbf{v}^\top\mathbf{E}^\top + \mathbf{E}\mathbf{v}\mathbf{u}^\top)}_{\mathbf{G}}.\end{aligned}$$

The eigenvalue decomposition of \mathbf{B} can be written as

$$\mathbf{B} = [\mathbf{u} \ \mathbf{U}_2] \begin{bmatrix} \beta^2 + m\sigma^2 & \\ & m\sigma^2\mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{u}^\top \\ \mathbf{U}_2^\top \end{bmatrix}.$$

We first show a deterministic lower bound for $|\langle \hat{\mathbf{u}}, \mathbf{u} \rangle|$ assuming $\beta^2 \geq 2\|\mathbf{G}\|$, where $\hat{\mathbf{u}}$ is the leading eigenvector of $\tilde{\mathbf{X}}\tilde{\mathbf{X}}^\top$. Then we bound the spectral norm $\|\mathbf{G}\|$ of the noise term (Lemma 3) and derive the sufficient condition for β .

Let $\hat{\mathbf{u}}$ be the leading eigenvector of $\tilde{\mathbf{X}}\tilde{\mathbf{X}}^\top$ with eigenvalue $\hat{\lambda}$, $\mathbf{r} = \mathbf{B}\hat{\mathbf{u}} - \hat{\lambda}\hat{\mathbf{u}} = -\mathbf{G}\hat{\mathbf{u}}$. We have $\mathbf{U}_2^\top\mathbf{r} = (m\sigma^2 - \hat{\lambda})\mathbf{U}_2^\top\hat{\mathbf{u}}$. Hence, for all $\beta^2 > 2\|\mathbf{G}\|$, it holds that

$$|\sin(\hat{\mathbf{u}}, \mathbf{u})| = \|\mathbf{U}_2^\top\hat{\mathbf{u}}\|_2 = \frac{\|\mathbf{U}_2^\top\mathbf{r}\|_2}{\hat{\lambda} - m\sigma^2} \leq \frac{\|\mathbf{G}\|}{\beta^2 - \|\mathbf{G}\|} \leq \frac{2\|\mathbf{G}\|}{\beta^2},$$

where we used $\|\mathbf{U}_2^\top\mathbf{r}\|_2 = \|\mathbf{U}_2^\top\mathbf{G}\hat{\mathbf{u}}\|_2 \leq \|\mathbf{G}\|$, and $\hat{\lambda} \geq \mathbf{u}^\top\tilde{\mathbf{X}}\tilde{\mathbf{X}}^\top\mathbf{u} \geq \beta^2 + m\sigma^2 - \|\mathbf{G}\|$. Therefore,

$$|\langle \hat{\mathbf{u}}, \mathbf{u} \rangle| = |\cos(\hat{\mathbf{u}}, \mathbf{u})| \geq \sqrt{1 - \frac{4\|\mathbf{G}\|^2}{\beta^4}} \geq 1 - \frac{4\|\mathbf{G}\|^2}{\beta^4},$$

if $\beta^2 \geq 2\|\mathbf{G}\|$.

It follows from Lemma 3 (shown below) that

$$\|\mathbf{G}\| \leq \begin{cases} 2\bar{C}\sigma^2\sqrt{mn}, & \text{if } \beta/\sigma < \sqrt{m}, \\ 2\bar{C}\beta\sigma\sqrt{n}, & \text{otherwise,} \end{cases}$$

where \bar{C} is a universal constant with probability at least $1 - 4e^{-n}$.

Now consider the first case ($\beta/\sigma < \sqrt{m}$) and assume $\beta^2 \geq 4\bar{C}\sigma^2\sqrt{mn} \geq 2\|\mathbf{G}\|$. Note that this case only arises when $\sqrt{m} \geq 4\bar{C}\sqrt{n}$. Denoting $C = 16\bar{C}^2$, we obtain the first case in the theorem. Next, consider the second case ($\beta/\sigma \geq \sqrt{m}$). If $\sqrt{m} \geq 4\bar{C}\sqrt{n}$ as above, we have $\beta/\sigma \geq 4\bar{C}\sqrt{n}$, which implies $\beta^2 \geq 2\|\mathbf{G}\|$ and we obtain the second case in the theorem. On the other hand, if $\sqrt{m} < 4\bar{C}\sqrt{n}$, we require $\beta/\sigma \geq 4\bar{C}\sqrt{n}$ to obtain the last case in the theorem.

Lemma 3. *Let \mathbf{G} be constructed as in Theorem 1. If $m \geq n$, there exists an universal constant \bar{C} such that*

$$\|\mathbf{G}\| \leq \bar{C}\sigma^2 \left(\sqrt{mn} + \sqrt{n(\beta/\sigma)^2} \right),$$

with probability at least $1 - 4e^{-n}$.

Proof. The proof is an ε -net argument. Let

$$\lambda = 2\sigma^2 \left(\sqrt{4mn} + 4n + \sqrt{8n(\beta/\sigma)^2} \right).$$

The goal is to control $|\mathbf{x}^\top\mathbf{G}\mathbf{x}|$ for all the vectors \mathbf{x} on the unit Euclidean sphere \mathcal{S}^{n-1} . In order to do this, we first bound the probability of the tail event $|\mathbf{x}^\top\mathbf{G}\mathbf{x}| > \lambda$, for any fixed $\mathbf{x} \in \mathcal{S}^{n-1}$. Then

we bound the probability that $|\mathbf{x}^\top \mathbf{G} \mathbf{x}| > \lambda$ for all the vectors in a ε -net \mathcal{N}_ε . Finally, we establish the connection between $\sup_{\mathbf{x} \in \mathcal{N}_\varepsilon} |\mathbf{x}^\top \mathbf{G} \mathbf{x}|$ and $\|\mathbf{G}\|$.

To bound $\mathbb{P}(|\mathbf{x}^\top \mathbf{G} \mathbf{x}| > \lambda)$ for a fix $\mathbf{x} \in \mathcal{S}^{n-1}$, we expand $\mathbf{x}^\top \mathbf{G} \mathbf{x}$ as

$$\mathbf{x}^\top \mathbf{G} \mathbf{x} = \sigma^2(\|\mathbf{z}\|^2 - m) + 2\beta\sigma(\mathbf{u}^\top \mathbf{x})\gamma,$$

where $\mathbf{z} = \mathbf{E}^\top \mathbf{x}$ and $\gamma = \mathbf{v}^\top \mathbf{z}$. Since $\mathbf{z} \sim \mathcal{N}(0, \mathbf{I})$, we can see that $\|\mathbf{z}\|^2$ is χ^2 distributed with m degrees of freedom and $\gamma \sim \mathcal{N}(0, 1)$.

First we bound the deviation of the χ^2 term. By the corollary of Lemma 1 in [17], we have

$$\mathbb{P}(|\|\mathbf{z}\|^2 - m| > \lambda_1) \leq 2e^{-4n}, \quad (13)$$

where $\lambda_1 = 2(\sqrt{4mn} + 4n)$.

Next we bound the deviation of the Gaussian term. Using the Gaussian tail inequality, we have

$$\mathbb{P}(|\gamma| > \lambda_2) \leq 2e^{-4n}, \quad (14)$$

where $\lambda_2 = \sqrt{8n}$.

Combining inequalities (22) and (14), we have

$$\begin{aligned} & \mathbb{P}(|\mathbf{x}^\top \mathbf{G} \mathbf{x}| > \lambda) \\ & \leq \mathbb{P}(\sigma^2|\|\mathbf{z}\|^2 - m| + |2\beta\sigma(\mathbf{u}^\top \mathbf{x})\gamma| > \sigma^2\lambda_1 + 2\beta\sigma\lambda_2) \\ & \leq \mathbb{P}(|\|\mathbf{z}\|^2 - m| > \lambda_1 \vee |\gamma| > \lambda_2) \\ & \leq \mathbb{P}(|\|\mathbf{z}\|^2 - m| > \lambda_1) + \mathbb{P}(|\gamma| > \lambda_2) \\ & \leq 4e^{-4n}, \end{aligned}$$

where the second to last line follows from the union bound.

Furthermore, using Lemma 5.2 and 5.4 of [27], for any $\varepsilon \in [0, 1)$, it holds that

$$|\mathcal{N}_\varepsilon| \leq (1 + 2/\varepsilon)^n,$$

and

$$\|\mathbf{G}\| \leq (1 - 2\varepsilon)^{-1} \sup_{\mathbf{x} \in \mathcal{N}_\varepsilon} |\mathbf{x}^\top \mathbf{G} \mathbf{x}|.$$

Taking the union bound over all the vectors in $\mathcal{N}_{1/4}$, we obtain

$$\mathbb{P}\left(\sup_{\mathbf{x} \in \mathcal{N}_{1/4}} |\mathbf{x}^\top \mathbf{G} \mathbf{x}| > \lambda\right) \leq |\mathcal{N}_{1/4}| 4e^{-4n} < 4e^{-n}.$$

Finally, the statement is obtained by noticing that $n \leq m$. \square

We prove a more general version of the theorem that allows the signal part to be rank R in Appendix E.

D.2 Proof of Lemma 1

Proof. By definition,

$$\begin{aligned} \|\mathcal{Y}\|_{s^*} &= \sup_{\{\mathbf{M}^{(k)}\}_{k=1}^K} \langle \mathcal{Y}, \sum_{k=1}^K \text{fold}_k(\mathbf{M}^{(k)} \mathbf{S}^{(k)\top}) \rangle \\ & \text{s.t. } \sum_{k=1}^K \|\mathbf{M}^{(k)}\|_* \leq 1 \\ &= \sup_{\{\mathbf{M}^{(k)}\}_{k=1}^K} \sum_{k=1}^K \langle \mathbf{Y}_{(k)} \mathbf{S}^{(k)}, \mathbf{M}^{(k)} \rangle \\ & \text{s.t. } \sum_{k=1}^K \|\mathbf{M}^{(k)}\|_* \leq 1 \\ &= \max_k \|\mathbf{Y}_{(k)} \mathbf{S}^{(k)}\|, \end{aligned}$$

where we used the Hölder inequality in the last line. \square

D.3 Proof of Theorem 2

First we decompose the error as

$$\|\mathcal{X}^* - \hat{\mathcal{X}}\|_F \leq \|\mathcal{X}^* - \mathcal{X}_p\|_F + \|\mathcal{X}_p - \hat{\mathcal{X}}\|_F.$$

The first term is an approximation error that depends on the choice of the subspace $\mathcal{S}^{(k)}$. The second term corresponds to an estimation error and we analyze the second term below.

Since $\hat{\mathcal{X}}$ is the minimizer of (10) and \mathcal{X}_p is feasible,

$$\frac{1}{2}\|\mathcal{Y} - \hat{\mathcal{X}}\|_F^2 + \lambda \sum_{k=1}^K \|\hat{\mathcal{M}}^{(k)}\|_* \leq \frac{1}{2}\|\mathcal{Y} - \mathcal{X}_p\|_F^2 + \lambda \sum_{k=1}^K \|\mathcal{M}_p^{(k)}\|_*,$$

from which we have

$$\frac{1}{2}\|\mathcal{X}_p - \hat{\mathcal{X}}\|_F^2 \leq \|\mathcal{Y} - \mathcal{X}_p\|_{s^*} \|\mathcal{X}_p - \hat{\mathcal{X}}\|_s + \lambda \sum_{k=1}^K \left(\|\mathcal{M}_p^{(k)}\|_* - \|\hat{\mathcal{M}}^{(k)}\|_* \right). \quad (15)$$

Next we define $\Delta_k := \hat{\mathcal{M}}^{(k)} - \mathcal{M}_p^{(k)} \in \mathbb{R}^{n_k \times H^{K-1}}$ and define its orthogonal decomposition $\Delta_k = \Delta'_k + \Delta''_k$ as

$$\Delta''_k := (\mathbf{I}_{n_k} - P_{U_p})\Delta_k(\mathbf{I}_{H^{K-1}} - P_{V_p}),$$

where P_{U_p} and P_{V_p} are projection matrices to the column and row spaces of $\mathcal{M}_p^{(k)}$, respectively, and $\Delta'_k := \Delta_k - \Delta''_k$.

The above definition allows us to decompose $\|\hat{\mathcal{M}}^{(k)}\|_*$ as follows:

$$\begin{aligned} \|\hat{\mathcal{M}}^{(k)}\|_* &= \|\mathcal{M}_p^{(k)} + \Delta''_k + \Delta'_k\|_* \\ &\geq \|\mathcal{M}_p^{(k)}\|_* + \|\Delta''_k\|_* - \|\Delta'_k\|_*. \end{aligned} \quad (16)$$

Moreover,

$$\|\mathcal{X}_p - \hat{\mathcal{X}}\|_s \leq \sum_{k=1}^K \|\Delta_k\|_* \leq \sum_{k=1}^K (\|\Delta'_k\|_* + \|\Delta''_k\|_*) \quad (17)$$

Combining inequalities (15)–(17), we have

$$\frac{1}{2}\|\mathcal{X}_p - \hat{\mathcal{X}}\|_F^2 \leq (\|\mathcal{Y} - \mathcal{X}_p\|_{s^*} + \lambda) \sum_{k=1}^K \|\Delta'_k\|_* + (\|\mathcal{Y} - \mathcal{X}_p\|_{s^*} - \lambda) \sum_{k=1}^K \|\Delta''_k\|_*. \quad (18)$$

Since

$$\|\mathcal{Y} - \mathcal{X}_p\|_{s^*} \leq \sigma \|\mathcal{E}\|_{s^*} + \|\mathcal{X}^* - \mathcal{X}_p\|_{s^*},$$

if $\lambda \geq \sigma \|\mathcal{E}\|_{s^*} + \|\mathcal{X}^* - \mathcal{X}_p\|_{s^*}$, the second term in the right-hand side of inequality (18) can be ignored and we have

$$\begin{aligned} \frac{1}{2}\|\mathcal{X}_p - \hat{\mathcal{X}}\|_F^2 &\leq 2\lambda \sum_{k=1}^K \|\Delta'_k\|_* \\ &\leq 2\lambda \sum_{k=1}^K \sqrt{2r_k} \|\Delta'_k\|_F \\ &\leq 2\lambda \sum_{k=1}^K \sqrt{2r_k} \|\Delta_k\|_F \\ &\leq 2\sqrt{2}\lambda \sqrt{\sum_{k=1}^K r_k} \sqrt{\sum_{k=1}^K \|\Delta_k\|_F^2}, \end{aligned} \quad (19)$$

where in the second line we used a simple observation that $\text{rank}(\Delta'_k) \leq 2r_k$.

Next, we relate the norm $\|\mathcal{X}_p - \hat{\mathcal{X}}\|_F$ to the sum $\sum_{k=1}^K \|\Delta_k\|_F^2$ in the right-hand side of inequality (19).

First suppose that $\sum_{k=1}^K \|\Delta_k\|_F^2 \leq \|\mathcal{X}_p - \hat{\mathcal{X}}\|_F^2$. Then from inequality (19), we have

$$\|\mathcal{X}_p - \hat{\mathcal{X}}\|_F \leq 4\sqrt{2}\lambda \sqrt{\sum_{k=1}^K r_k}$$

by dividing both sides by $\|\mathcal{X}_p - \hat{\mathcal{X}}\|_F$.

On the other hand, if $\|\mathcal{X}_p - \hat{\mathcal{X}}\|_F^2 \leq \sum_{k=1}^K \|\Delta_k\|_F^2$, we use the following lemma

Lemma 4. *Suppose $\{\mathbf{M}_p^{(k)}\}_{k=1}^K, \{\hat{\mathbf{M}}^{(k)}\}_{k=1}^K \in \mathcal{M}(\rho)$, and $\mathbf{S}^{(k)}$ is constructed as a Kronecker product of $K-1$ ortho-normal matrices $\hat{\mathbf{P}}^{(\ell)}$ as $\mathbf{S}^{(k)} = \hat{\mathbf{P}}^{(k-1)} \otimes \dots \otimes \hat{\mathbf{P}}^{(k+1)}$, where $(\hat{\mathbf{P}}^{(\ell)})^\top \hat{\mathbf{P}}^{(\ell)} = \mathbf{I}_H$ for $\ell = 1, \dots, K$. Then for $\mathcal{X}_p = \sum_{k=1}^K \text{fold}_k(\mathbf{M}_p^{(k)} \mathbf{S}^{(k)\top})$ and $\hat{\mathcal{X}} = \sum_{k=1}^K \text{fold}_k(\hat{\mathbf{M}}^{(k)} \mathbf{S}^{(k)\top})$, the following inequality holds:*

$$\frac{1}{2} \sum_{k=1}^K \|\Delta_k\|_F^2 \leq \frac{1}{2} \|\mathcal{X}_p - \hat{\mathcal{X}}\|_F^2 + \rho \max_k(\sqrt{n_k} + \sqrt{H^{K-1}}) \sum_{k=1}^K \|\Delta_k\|_*. \quad (20)$$

Proof. The proof is presented in Section D.5. □

Combining inequalities (18) and (20), we have

$$\begin{aligned} \frac{1}{2} \sum_{k=1}^K \|\Delta_k\|_F^2 &\leq \left(\|\mathcal{Y} - \mathcal{X}_p\|_{s^*} + \rho \max_k(\sqrt{n_k} + \sqrt{H^{K-1}}) + \lambda \right) \sum_{k=1}^K \|\Delta'_k\|_* \\ &\quad + \left(\|\mathcal{Y} - \mathcal{X}_p\|_{s^*} + \rho \max_k(\sqrt{n_k} + \sqrt{H^{K-1}}) - \lambda \right) \sum_{k=1}^K \|\Delta''_k\|_*. \end{aligned}$$

Thus if we take $\lambda \geq \sigma \|\mathcal{E}\|_{s^*} + \|\mathcal{X}^* - \mathcal{X}_p\|_{s^*} + \rho \max_k(\sqrt{n_k} + \sqrt{H^{K-1}})$, the second term in the right-hand side can be ignored and following the derivation leading to inequality (19) and dividing both sides by $\sqrt{\sum_{k=1}^K \|\Delta_k\|_F^2}$, we have

$$\|\mathcal{X}_p - \hat{\mathcal{X}}\|_F \leq \sqrt{\sum_{k=1}^K \|\Delta_k\|_F^2} \leq 4\sqrt{2}\lambda \sqrt{\sum_{k=1}^K r_k},$$

where the first inequality follows from the assumption.

The final step of the proof is to bound the norm $\|\mathcal{E}\|_{s^*}$ with sufficiently high probability. By Lemma 1,

$$\|\mathcal{E}\|_{s^*} = \max_k \|\mathbf{E}_{(k)} \mathbf{S}^{(k)}\|.$$

Therefore, taking the union bound, we have

$$\mathbb{P} \left(\max_k \|\mathbf{E}_{(k)} \mathbf{S}^{(k)}\| \geq t \right) \leq \sum_{k=1}^K \mathbb{P} \left(\|\mathbf{E}_{(k)} \mathbf{S}^{(k)}\| \geq t \right). \quad (21)$$

Now since each $\mathbf{E}_{(k)} \mathbf{S}^{(k)} \in \mathbb{R}^{n_k \times H^{K-1}}$ is a random matrix with i.i.d. standard Gaussian entries,

$$\mathbb{P} \left(\|\mathbf{E}_{(k)} \mathbf{S}^{(k)}\| \geq \sqrt{n_k} + \sqrt{H^{K-1}} + t \right) \leq \exp(-t^2/(2\sigma^2)).$$

Therefore, choosing $t = \max_k(\sqrt{n_k} + \sqrt{H^{K-1}}) + \sqrt{2\log(K/\delta)}$ in inequality (21), we have

$$\max_k \|\mathbf{E}_{(k)} \mathbf{S}^{(k)}\| \leq \max_k(\sqrt{n_k} + \sqrt{H^{K-1}}) + \sqrt{2\log(K/\delta)},$$

with probability at least $1 - \delta$. Plugging this into the condition for the regularization parameter λ , we obtain what we wanted.

D.4 Proof of Lemma 2

Proof. i) Let $\otimes_{k' \in [K] \setminus k} \mathbf{U}^{(k')}$ denote $\mathbf{U}^{(1)} \otimes \dots \otimes \mathbf{U}^{(k-1)} \otimes \mathbf{U}^{(k+1)} \otimes \dots \otimes \mathbf{U}^{(K)}$. We have

$$\begin{aligned} \mathbf{X}_{(k)}^* &= \mathbf{U}^{(k)} \mathbf{C}_{(k)} \left(\otimes_{k' \in [K] \setminus k} \mathbf{U}^{(k')} \right)^\top \\ &= \mathbf{U}^{(k)} \mathbf{C}_{(k)} \left(\otimes_{k' \in [K] \setminus k} (\mathbf{U}^{(k')})^\top \right) \\ &= \mathbf{P}^{(k)} \mathbf{\Lambda}^{(k)} (\mathbf{Q}^{(k)})^\top. \end{aligned}$$

Because of the minimality of the Tucker decomposition (5), $\mathbf{X}_{(k)}^*$, $\mathbf{C}_{(k)}$ and $\mathbf{U}^{(k)}$ are all of rank R , for all $k \in [K]$. Therefore, both $\mathbf{C}_{(k)}$ and $\otimes_{k' \in [K] \setminus k} (\mathbf{U}^{(k')})^\top$ have full row rank.

Hence, $\mathbf{C}_{(k)}$ has a Moore-Penrose pseudo inverse $\mathbf{C}_{(k)}^\dagger$ such that $\mathbf{C}_{(k)} \mathbf{C}_{(k)}^\dagger = \mathbf{I}$, and so does $\otimes_{k' \in [K] \setminus k} (\mathbf{U}^{(k')})^\top$. As a result, we have

$$\mathbf{U}^{(k)} = \mathbf{P}^{(k)} \mathbf{\Lambda}^{(k)} (\mathbf{Q}^{(k)})^\top \left(\otimes_{k' \in [K] \setminus k} (\mathbf{U}^{(k')})^\top \right)^\dagger \mathbf{C}_{(k)}^\dagger.$$

ii) Similarly, we have

$$\mathbf{Q}^{(k)} \mathbf{\Lambda}^{(k)} (\mathbf{P}^{(k)})^\top = \left(\otimes_{k' \in [K] \setminus k} \mathbf{U}^{(k')} \right) \mathbf{C}_{(k)}^\top (\mathbf{U}^{(k)})^\top.$$

By the definition of SVD, $\mathbf{\Lambda}$ is invertible and $(\mathbf{P}^{(k)})^\top \mathbf{P}^{(k)} = \mathbf{I}$. Hence,

$$\mathbf{Q}^{(k)} = \left(\otimes_{k' \in [K] \setminus k} \mathbf{U}^{(k')} \right) \mathbf{C}_{(k)}^\top (\mathbf{U}^{(k)})^\top \mathbf{P}^{(k)} (\mathbf{\Lambda}^{(k)})^{-1}.$$

This means $\mathbf{Q}^{(k)} \in \text{span} \left(\otimes_{k' \in [K] \setminus k} \mathbf{U}^{(k')} \right)$ and we then conclude $\mathbf{Q}^{(k)} \in \text{span} \left(\otimes_{k' \in [K] \setminus k} \mathbf{P}^{(k')} \right)$ using (i). □

D.5 Proof of Lemma 4

Expanding \mathcal{X}_p and $\hat{\mathcal{X}}$, we have

$$\begin{aligned}
\|\mathcal{X}_p - \hat{\mathcal{X}}\|_F^2 &= \left\| \sum_{k=1}^K \text{fold}_k(\Delta_k \mathbf{S}^{(k)\top}) \right\|_F^2 \\
&= \sum_{k=1}^K \|\Delta_k\|_F^2 + \sum_{k \neq \ell} \langle \text{fold}_k(\Delta_k \mathbf{S}^{(k)\top}), \text{fold}_\ell(\Delta_\ell \mathbf{S}^{(\ell)\top}) \rangle \\
&= \sum_{k=1}^K \|\Delta_k\|_F^2 + \sum_{k \neq \ell} \langle \text{fold}_k(\Delta_k) \times_{k' \neq k} \hat{\mathbf{P}}^{(k')}, \text{fold}_\ell(\Delta_\ell) \times_{\ell' \neq \ell} \hat{\mathbf{P}}^{(\ell')} \rangle \\
&= \sum_{k=1}^K \|\Delta_k\|_F^2 + \sum_{k \neq \ell} \langle \text{fold}_k(\Delta_k) \times_\ell \hat{\mathbf{P}}^{(\ell)}, \text{fold}_\ell(\Delta_\ell) \times_k \hat{\mathbf{P}}^{(k)} \rangle \\
&= \sum_{k=1}^K \|\Delta_k\|_F^2 - \sum_{k \neq \ell} \langle \Delta_k (\mathbf{I}_H \otimes \cdots \otimes \hat{\mathbf{P}}^{(\ell)} \otimes \cdots \otimes \mathbf{I}_H)^\top, \hat{\mathbf{P}}^{(k)} (\text{fold}_\ell(\Delta_\ell))_{(k)} \rangle \\
&\geq \sum_{k=1}^K \|\Delta_k\|_F^2 - \sum_{k \neq \ell} \|\Delta_k\|_* \cdot \|(\text{fold}_\ell(\Delta_\ell))_{(k)}\| \\
&\geq \sum_{k=1}^K \|\Delta_k\|_F^2 - 2\rho \max_k (\sqrt{n_k} + \sqrt{H^{K-1}}) \sum_{k=1}^K \|\Delta_k\|_*,
\end{aligned}$$

from which the lemma holds. Here we regarded $\text{fold}_k(\Delta_k \mathbf{S}^{(k)})$ as a Tucker decomposition with the core tensor $\text{fold}_k(\Delta_k)$ and factor matrices $\hat{\mathbf{P}}^{(k')}$ for $k' \neq k$. Most of the factors except for k and ℓ cancel out when calculating the inner product between two such tensors in the third line, because $(\hat{\mathbf{P}}^{(k')})^\top \hat{\mathbf{P}}^{(k')} = \mathbf{I}_H$. After unfolding the inner product at the k th mode in the fifth line, we notice that a multiplication by an ortho-normal matrix does not affect the nuclear norm or the spectral norm. In the last line we used $\{\Delta_k\}_{k=1}^K \in \mathcal{M}(2\rho)$, which follows from the assumption that both $\{\mathbf{M}_p^{(k)}\}_{k=1}^K, \{\hat{\mathbf{M}}^{(k)}\}_{k=1}^K \in \mathcal{M}(\rho)$. \square

E Generalization of Theorem 1 to the higher rank case

Theorem 4. Suppose that $\mathbf{X} = \sum_{r=1}^R \beta_r \mathbf{u}_r \mathbf{v}_r^\top$, where $\mathbf{u}_1, \dots, \mathbf{u}_R \in \mathbb{R}^n$ and $\mathbf{v}_1, \dots, \mathbf{v}_R \in \mathbb{R}^m$ are unit orthogonal vectors respectively. Let $\tilde{\mathbf{X}} = \mathbf{X} + \sigma \mathbf{E}$ be the noisy observation of \mathbf{X} . There exists an universal constant C such that with probability at least $1 - 3e^{-n}$, if $m/n \geq C(\beta_1/\beta_R)^4$, then

$$|\cos(\hat{\mathbf{U}}, \mathbf{U})| \geq \begin{cases} 1 - \frac{Cmn}{(\beta_R/\sigma)^4}, & \text{if } \frac{\beta_1}{\sqrt{m}} < \sigma \leq \frac{\beta_R}{(Cmn)^{\frac{1}{4}}}, \\ 1 - \frac{Cn(\beta_1/\beta_R)^2}{(\beta_R/\sigma)^2}, & \text{if } \sigma \leq \frac{\beta_1}{\sqrt{m}}, \end{cases}$$

otherwise, $|\cos(\hat{\mathbf{U}}, \mathbf{U})| \geq 1 - \frac{Cn(\beta_1/\beta_R)^2}{(\beta_R/\sigma)^2}$ if $\sigma \leq \beta_R^2/(Cn)^{\frac{1}{2}}\beta_1$.

Suppose that $\mathbf{X} = \sum_{r=1}^R \beta_r^2 \mathbf{u}_r \mathbf{v}_r^\top$ and $\tilde{\mathbf{X}} = \mathbf{X} + \sigma \mathbf{E}$. We consider the second moment of $\tilde{\mathbf{X}}$:

$$\begin{aligned} \tilde{\mathbf{X}} \tilde{\mathbf{X}}^\top &= \sum_{r=1}^R \beta_r^2 \mathbf{u}_r \mathbf{u}_r^\top + \sigma \left(\sum_{r=1}^R \beta_r (\mathbf{u}_r \mathbf{v}_r^\top \mathbf{E}^\top + \mathbf{E} \mathbf{v}_r \mathbf{u}_r^\top) \right) + \sigma^2 \mathbf{E} \mathbf{E}^\top \\ &= \underbrace{\sum_{r=1}^R \beta_r^2 \mathbf{u}_r \mathbf{u}_r^\top + m\sigma^2 \mathbf{I}}_{\mathbf{B}} + \underbrace{\sigma^2 \mathbf{E} \mathbf{E}^\top - m\sigma^2 \mathbf{I} + \sigma \left(\sum_{i=1}^R \beta_r (\mathbf{u}_r \mathbf{v}_r^\top \mathbf{E}^\top + \mathbf{E} \mathbf{v}_r \mathbf{u}_r^\top) \right)}_{\mathbf{G}}. \end{aligned}$$

The eigenvalue decomposition of \mathbf{B} can be written as

$$\mathbf{B} = [\mathbf{U} \ \mathbf{U}_2] \begin{bmatrix} \Sigma + m\sigma^2 \mathbf{I} & \\ & m\sigma^2 \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{U}^\top \\ \mathbf{U}_2^\top \end{bmatrix},$$

where $\mathbf{U} \in \mathbb{R}^{n \times R}$ and $\Sigma = \text{diag}(\beta_1^2, \dots, \beta_R^2)$. Similarly, the eigenvalue decomposition of $\tilde{\mathbf{X}} \tilde{\mathbf{X}}^\top$ can be written as

$$\tilde{\mathbf{X}} \tilde{\mathbf{X}}^\top = [\hat{\mathbf{U}} \ \hat{\mathbf{U}}_2] \begin{bmatrix} \hat{\Sigma} & \\ & \hat{\Sigma}' \end{bmatrix} \begin{bmatrix} \hat{\mathbf{U}}^\top \\ \hat{\mathbf{U}}_2^\top \end{bmatrix},$$

where $\hat{\Sigma} = \text{diag}(\hat{\lambda}_1, \dots, \hat{\lambda}_R)$ and $\hat{\Sigma}' = \text{diag}(\hat{\lambda}_{R+1}, \dots, \hat{\lambda}_n)$ s.t. $\hat{\lambda}_1 \geq \dots \geq \hat{\lambda}_n$ are the eigenvalues of $\tilde{\mathbf{X}} \tilde{\mathbf{X}}^\top$.

We first show a deterministic lower bound for $|\sin(\hat{\mathbf{U}}, \mathbf{U})|$ assuming $\beta_R^2 \geq 2\|\mathbf{G}\|$. Then we bound the spectral norm $\|\mathbf{G}\|$ of the noise term (Lemma 3) and derive the sufficient condition for β_R^2 .

The maximum singular value of $m\sigma^2 \mathbf{I}$ is $m\sigma^2$. The minimum singular value of $\hat{\Sigma}$ is $|\hat{\lambda}_R|$. By Wely's theorem, $\|\mathbf{G}\| \geq |\hat{\lambda}_R - \beta_R^2 - m\sigma^2|$, which means

$$\hat{\lambda}_R \geq m\sigma^2 + \beta_R^2 - \|\mathbf{G}\|.$$

Let $\mathbf{R} = \mathbf{G}\hat{\mathbf{U}}$. Since $\beta_R^2 \geq 2\|\mathbf{G}\|$, we can apply the Wedin theorem and obtain

$$|\sin(\hat{\mathbf{U}}, \mathbf{U})| = \|\mathbf{U}_2^\top \hat{\mathbf{U}}\| \leq \frac{\|\mathbf{R}\|}{\beta_R^2 - \|\mathbf{G}\|} = \frac{\|\mathbf{G}\hat{\mathbf{U}}\|}{\beta_R^2 - \|\mathbf{G}\|} \leq \frac{\|\mathbf{G}\|}{\beta_R^2 - \|\mathbf{G}\|} \leq \frac{2\|\mathbf{G}\|}{\beta_R^2},$$

where we used the property that the spectral norm is sub-multiplicative and $\|\hat{\mathbf{U}}\| = 1$ in the second to last step.

Therefore,

$$|\cos(\hat{\mathbf{U}}, \mathbf{U})| \geq \sqrt{1 - \frac{4\|\mathbf{G}\|^2}{\beta_R^4}} \geq 1 - \frac{4\|\mathbf{G}\|^2}{\beta_R^4},$$

if $\beta_R^2 \geq 2\|\mathbf{G}\|$. It follows from Lemma 5 (shown below) that with probability at least $1 - 3e^{-n}$

$$\|\mathbf{G}\| \leq \begin{cases} 2\bar{C}\sigma^2\sqrt{mn}, & \text{if } \beta_1/\sigma < \sqrt{m}, \\ 2\bar{C}\sigma\sqrt{n}\beta_1, & \text{otherwise,} \end{cases}$$

where \bar{C} is an universal constant.

Let $C = 16\bar{C}^2$. Now consider the first situation where $m/n > C(\beta_1/\beta_R)^4$. If $\sigma > \frac{\beta_1}{\sqrt{m}}$, we have $\|\mathbf{G}\| \leq \frac{\sigma^2}{2}(Cmn)^{\frac{1}{2}}$. Meanwhile, if $\sigma \leq \frac{\beta_R}{(Cmn)^{\frac{1}{4}}}$, then we have $\beta_R^2 \geq \sigma^2(Cmn)^{\frac{1}{2}} \geq 2\|\mathbf{G}\|$.

Combining these two conditions we obtain the first case in the theorem. When $\sigma \leq \frac{\beta_1}{\sqrt{m}}$, we can see that $\|\mathbf{G}\| \leq \frac{\sigma}{2}(Cn)^{\frac{1}{2}}\beta_1$. Moreover, since $m/n > C(\beta_1/\beta_R)^4$, it is implied that $\sigma \leq \beta_R^2/(Cn)^{\frac{1}{2}}\beta_1$ and thus $\beta_R^2 \geq 2\|\mathbf{G}\|$. This gives us the second case.

On the other hand, if $m/n \leq C(\beta_1/\beta_R)^4$, we require $\sigma \leq \beta_R^2/(Cn)^{\frac{1}{2}}\beta_1$ to obtain the last case in the theorem.

Lemma 5. *Let \mathbf{G} be constructed as in the proof of Theorem 4. If $m \geq n$, there exists an universal constant \bar{C} such that*

$$\|\mathbf{G}\| \leq \bar{C}\sigma^2(\sqrt{mn} + \sqrt{n}\beta_1/\sigma),$$

with probability at least $1 - 3e^{-n}$.

Proof. The proof is an ε -net argument. Let

$$\lambda = 2\sigma^2 \left(\sqrt{4mn} + 4n + \sqrt{R + 8n + 4\sqrt{Rn}} \cdot \beta_1/\sigma \right).$$

The goal is to control $|\mathbf{x}^\top \mathbf{G}\mathbf{x}|$ for all the vectors \mathbf{x} on the unit Euclidean sphere \mathcal{S}^{n-1} . In order to do this, we first bound the probability of the tail event $|\mathbf{x}^\top \mathbf{G}\mathbf{x}| > \lambda$, for any fixed $\mathbf{x} \in \mathcal{S}^{n-1}$. Then we bound the probability that $|\mathbf{x}^\top \mathbf{G}\mathbf{x}| > \lambda$ for all the vectors in a ε -net \mathcal{N}_ε . Finally, we establish the connection between $\sup_{\mathbf{x} \in \mathcal{N}_\varepsilon} |\mathbf{x}^\top \mathbf{G}\mathbf{x}|$ and $\|\mathbf{G}\|$.

To bound $\mathbb{P}(|\mathbf{x}^\top \mathbf{G}\mathbf{x}| > \lambda)$ for a fix $\mathbf{x} \in \mathcal{S}^{n-1}$, we expand $\mathbf{x}^\top \mathbf{G}\mathbf{x}$ as

$$\mathbf{x}^\top \mathbf{G}\mathbf{x} = \sigma^2(\|\mathbf{z}\|^2 - m) + 2\sigma \sum_{r=1}^R \beta_r \gamma_r (\mathbf{u}_r^\top \mathbf{x}),$$

where $\mathbf{z} = \mathbf{E}^\top \mathbf{x}$ and $\gamma_r = \mathbf{v}_r^\top \mathbf{z}$. It is easy to see that $\gamma_r \sim \mathcal{N}(0, 1)$, $\mathbf{z} \sim \mathcal{N}(0, \mathbf{I})$ and $\|\mathbf{z}\|^2$ is χ^2 distributed with m degrees of freedom.

Let $\boldsymbol{\gamma} = [\gamma_1, \dots, \gamma_R]$ and $\boldsymbol{\omega} = [\mathbf{u}_1^\top \mathbf{x}, \dots, \mathbf{u}_R^\top \mathbf{x}]$. We have

$$\begin{aligned} |\mathbf{x}^\top \mathbf{G}\mathbf{x}| &\leq \sigma^2 \left| \|\mathbf{z}\|^2 - m \right| + 2\sigma \left| \sum_{r=1}^R \beta_r \gamma_r (\mathbf{u}_r^\top \mathbf{x}) \right| \\ &\leq \sigma^2 \left| \|\mathbf{z}\|^2 - m \right| + 2\sigma \sum_{r=1}^R \max_{r \in [R]} |\beta_r| \cdot |\gamma_r| \cdot |\mathbf{u}_r^\top \mathbf{x}| \\ &\leq \sigma^2 \left| \|\mathbf{z}\|^2 - m \right| + 2\sigma \beta_1 \cdot \|\boldsymbol{\gamma}\| \cdot \|\boldsymbol{\omega}\| \\ &\leq \sigma^2 \left| \|\mathbf{z}\|^2 - m \right| + 2\sigma \beta_1 \cdot \|\boldsymbol{\gamma}\|, \end{aligned}$$

where we used the Cauchy-Schwarz inequality in the second to last line and the fact $\|\boldsymbol{\omega}\| \leq 1$ in the last line. Note that $\gamma_1, \dots, \gamma_R$ are i.i.d standard Gaussian distributed so that $\|\boldsymbol{\gamma}\|^2$ is χ^2 distributed with R degrees.

First we bound the deviation of the χ_m^2 term. By the corollary of Lemma 1 in [17], we have

$$\mathbb{P}(\left| \|\mathbf{z}\|^2 - m \right| > \lambda_1) \leq 2e^{-4n}, \quad (22)$$

where $\lambda_1 = 2(\sqrt{4mn} + 4n)$.

Next we bound the χ_R^2 term. Similarly, we have

$$\mathbb{P}(\|\boldsymbol{\gamma}\|^2 - R > \lambda_2) \leq e^{-4n}, \quad (23)$$

where $\lambda_2 = 2(\sqrt{4Rn} + 4n)$.

Combining inequalities (22) and (23), we have

$$\begin{aligned} & \mathbb{P}(|\mathbf{x}^\top \mathbf{G} \mathbf{x}| > \lambda) \\ & \leq \mathbb{P}\left(\sigma^2 \left| \|\mathbf{z}\|^2 - m \right| + 2\sigma\beta_1 \|\boldsymbol{\gamma}\| > \sigma^2 \lambda_1 + 2\sigma\beta_1 \sqrt{R + \lambda_2}\right) \\ & \leq \mathbb{P}\left(\left| \|\mathbf{z}\|^2 - m \right| > \lambda_1 \vee \|\boldsymbol{\gamma}\| > \sqrt{R + \lambda_2}\right) \\ & \leq \mathbb{P}\left(\left| \|\mathbf{z}\|^2 - m \right| > \lambda_1\right) + \mathbb{P}\left(\|\boldsymbol{\gamma}\| > \sqrt{R + \lambda_2}\right) \\ & \leq 3e^{-4n}, \end{aligned}$$

where the second to last line follows from the union bound.

Furthermore, using Lemma 5.2 and 5.4 of [27], for any $\varepsilon \in [0, 1)$, it holds that

$$|\mathcal{N}_\varepsilon| \leq (1 + 2/\varepsilon)^n,$$

and

$$\|\mathbf{G}\| \leq (1 - 2\varepsilon)^{-1} \sup_{\mathbf{x} \in \mathcal{N}_\varepsilon} |\mathbf{x}^\top \mathbf{G} \mathbf{x}|.$$

Taking the union bound over all the vectors in $\mathcal{N}_{1/4}$, we obtain

$$\mathbb{P}(\|\mathbf{G}\| \leq 2\lambda) \leq \mathbb{P}\left(\sup_{\mathbf{x} \in \mathcal{N}_{1/4}} |\mathbf{x}^\top \mathbf{G} \mathbf{x}| > \lambda\right) \leq |\mathcal{N}_{1/4}| 3e^{-4n} < 3e^{-n}.$$

Finally, the statement is obtained by noticing that $n \leq m$ and $R \leq n$. □