On the Error-Propagation of Inexact Hotelling's Deflation for Principal Component Analysis

Fangshuo Liao[†], Junhyung Lyle Kim[†], Cruz Barnum[*], Anastasios Kyrillidis[†]
[†]Rice University Computer Science Department, {Fangshuo.Liao,jlylekim,anastasios}@rice.edu
[*]UIUC Computer Science Department, cruzjb2@illinois.edu

Background: PCA and Deflation Method

Our focus: Real symmetric matrix $\mathbf{\Sigma} \in \mathbb{R}^{d\times d}$ with eigenvalues and eigenvectors $\{\lambda_i^*, \mathbf{u}_i^*\}_{i=1}^d$, satisfying $1 = \lambda_1^* > \lambda_2^* > \cdots > \lambda_d^* > 0$.

Target: Find the top-K eigenvectors of $\mathbf{\Sigma}$. When $K = 1$, one can solve the following problem using e.g. power iteration.

$$\mathbf{u}^* = \operatorname*{arg\,min}_{\mathbf{v}\in\mathbb{R}^d, \|\mathbf{v}\|_2=1} \mathbf{v}^\top \mathbf{\Sigma} \mathbf{v}$$

When $K > 1$, the problem can be generalized to

$$\mathbf{U}^* = \operatorname*{arg\,max}_{\mathbf{V}\in\{\mathbf{Q}_{:,:K}:\, \mathbf{Q}\in SO(d)\}} \langle \mathbf{\Sigma}\mathbf{V}, \mathbf{V}\rangle, \tag{1}$$

A popular method to solve (1) is the deflation method

$$\mathbf{\Sigma}_1 = \mathbf{\Sigma}; \quad \mathbf{v}_k = \mathtt{PCA}(\mathbf{\Sigma}_k, t); \qquad \mathbf{\Sigma}_{k+1} = \mathbf{\Sigma}_k - \mathbf{v}_k\mathbf{v}_k^\top\mathbf{\Sigma}_k\mathbf{v}_k\mathbf{v}_k^\top, \tag{2}$$

Here $\mathtt{PCA}(\cdot)$ denotes a sub-routine that solves for the top eigenvector of a matrix (e.g. power iteration) and t is the number of iterations of the sub-routine.

Ideal Scenario: Let $\mathbf{u}_k$ be normalized top eigenvector of $\mathbf{\Sigma}_k$ such that $\mathbf{u}_k^\top \mathbf{u}_k^* = 1$. When $\mathbf{v}_k = \mathbf{u}_k$ for all $k = 1, \ldots, K$, we can further guarantee that $\mathbf{u}_k = \mathbf{u}_k^*$.

Issue with Inexact Deflation

Issue: $\mathtt{PCA}(\cdot)$ is usually inexact: $\boldsymbol{\delta}_k := \mathbf{v}_k - \mathbf{u}_k \neq 0$.

Error Propagation: Let $\mathbf{\Sigma}_k^* = \sum_{k'=k+1} \lambda_k'^* \mathbf{u}_{k'}^*\mathbf{u}_{k'}^{*\top}$:

$$\mathbf{v}_k \neq \mathbf{u}_k \Rightarrow \mathbf{v}_k \neq \mathbf{u}_k^* \Rightarrow \mathbf{\Sigma}_{k+1} \neq \mathbf{\Sigma}_{k+1}^* \Rightarrow \mathbf{u}_{k+1} \neq \mathbf{u}_{k+1}^*$$

Influence on downstream task: $\|\boldsymbol{\delta}_k\|_2$ is larger when t is smaller (less computation).

Spectral Clustering on MNIST dataset

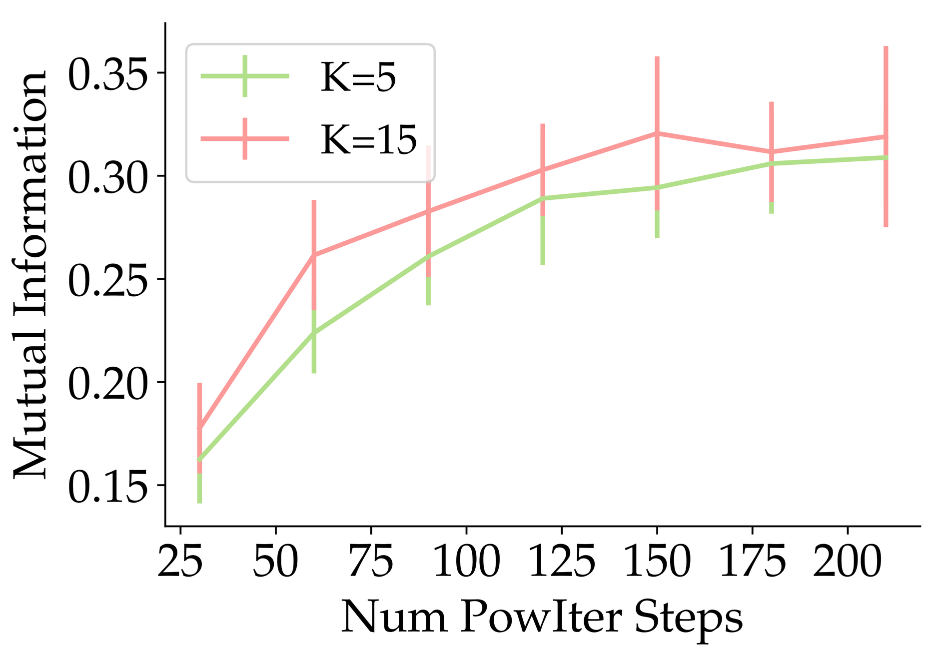

Central Question:

How can we characterize the accumulation and propagation of the errors from inexactly solving each top eigenvector in $\mathtt{PCA}(\cdot)$?

Sub-routine Agnostic Error Propagation

In this section, we assume that the sub-routines $\mathtt{PCA}(\cdot)$ only gives us the magnitude of the error $\|\boldsymbol{\delta}_k\|_2$'s.

Theorem 1. *Consider the scenario of looking for the top-K eigenvectors of $\mathbf{\Sigma} \in \mathbb{R}^{d\times d}$. Let $\mathcal{T}_j = \lambda_j^* - \lambda_{j+1}^*$, for $j \in [d-1]$, and $\mathcal{T}_d = \lambda_d^*$. Also, let $\mathcal{T}_{K,\min} = \min_{k\in[K]} \mathcal{T}_k$. If $\forall k \in [K-1]$ it holds:*[a]

$$\|\boldsymbol{\delta}_k\|_2 \leq \frac{\mathcal{T}_{K,\min}}{20K} \prod_{j=k+1}^{K-1} \left(3 + \frac{2\lambda_j^*}{\mathcal{T}_j}\right)^{-1} \Rightarrow; \|\mathbf{\Sigma}_k - \mathbf{\Sigma}_k^*\|_F \leq O(\mathcal{T}_k) \tag{3}$$

then the output of the deflation method in (2) satisfies for all $k \in [K]$:

$$\|\mathbf{v}_k - \mathbf{u}_k^*\|_2 \leq \|\mathbf{u}_k - \mathbf{u}_k^*\|_2 + \|\boldsymbol{\delta}_k\|_2 \leq 5 \sum_{k'=1}^{k} \frac{\lambda_{k'}^*}{\lambda_k^*} \|\boldsymbol{\delta}_{k'}\|_2 \prod_{j=k'+1}^{k} \left(3 + \frac{2\lambda_j^*}{\mathcal{T}_j}\right). \tag{4}$$

Example: Consider $\mathbf{\Sigma}$ with $\lambda_j^* = \frac{1}{j}$, and $\|\boldsymbol{\delta}_k\|_2 \leq c \cdot \alpha^t$ for some $\alpha \in (0, 1)$. Then

$$t \geq \Omega\left(\log\frac{1}{\epsilon} + K\log K\right) \Rightarrow \|\mathbf{v}_k - \mathbf{u}_k^*\|_2 \leq \epsilon;\ \forall k \in [K]$$

Issue: To guarantee an ϵ-small error, t grows faster with K as λ_j's decrease faster (smaller eigengaps).

[a] This theorem holds under a more relaxed assumption. The choice of $\|\boldsymbol{\delta}_k\|_2$ here is made just to make the result more interpretable.

Error Propagation When Using Power Iteration

In this section, we assume that the sub-routine $\mathtt{PCA}(\cdot)$ is the power iteration.

Power Iteration: Given a symmetric matrix $\mathbf{M} \in \mathbb{R}^{d\times d}$ with rank r and an initialization vector $\mathbf{x}_0 \in \mathbb{R}^d$, then power iteration executes:

$$\hat{\mathbf{x}}_{t+1} = \mathbf{M}\mathbf{x}_t; \quad \mathbf{x}_{t+1} = \frac{\hat{\mathbf{x}}_{t+1}}{\|\hat{\mathbf{x}}_{t+1}\|_2}. \tag{5}$$

Why is this helpful? Previous theorem only utilize the magnitude of $\boldsymbol{\delta}_k$'s, while power iteration also tells us the directional information of $\boldsymbol{\delta}_k$'s.

Theorem 2. *Consider the problem of looking for the top-K eigenvectors of $\mathbf{\Sigma} \in \mathbb{R}^{d\times d}$. Let $\mathcal{T}_i = \lambda_i^* - \lambda_{i+1}^*$ for $i \in [d-1]$ and $\mathcal{T}_d = \lambda_d^*$, and let $\mathcal{T}_{k,\min} = \min_{i\in[k]} \mathcal{T}_i$. Assume that there exists a constant $c_0 \in (0, \infty)$ such that the initialized vector $\mathbf{x}_{0,k}$ in the k-th power iteration procedure satisfies $\left|\mathbf{x}_{0,k}^\top \mathbf{u}_k\right| > c_0^{-1}$ for all $k \in [K]$.*[a] *Let $\mathcal{G} = \max_{k\in[K]}\left(1 + \frac{c_0\lambda_k^*\lambda_{k+1}^*}{\lambda_k^* - \lambda_{k+1}^*}\right)$. If we use t steps of power iteration to solve for the top eigenvectors such that:*

$$t \geq \begin{cases} \Omega\left(\frac{K - k + \log K/\mathcal{T}_{K,\min} + \log \mathcal{G}}{\log\left(7\lambda_k^* + \lambda_{k+1}^*\right) - \log\left(7\lambda_{k+1}^* + \lambda_k^*\right)}\right), & \forall k \leq K \\ \Omega\left(\frac{1}{\log \lambda_k^* - \log \lambda_{k+1}^*}\right), & \forall k \leq d \end{cases} \tag{6}$$

Then we can guarantee that for all $k \in [K]$

$$\|\mathbf{v}_k - \mathbf{u}_k^*\|_2 \leq 3 \sum_{k'=1}^{k} 8^{k-k'} \frac{\lambda_{k'}^*}{\lambda_k^*}\left(5\|\boldsymbol{\delta}_{k'}\|_2 + \frac{7c_0}{\mathcal{T}_{k'}}\left(\frac{\lambda_{k'+1}^*}{\lambda_{k'}^*}\right)^t\right). \tag{7}$$

Example: Consider $\mathbf{\Sigma}$ with $\lambda_j^* = \frac{1}{j}$. Then $\mathcal{G} = O(1)$, and

$$t \geq \Omega\left(\log\frac{1}{\epsilon} + K\right) \Rightarrow \|\mathbf{v}_k - \mathbf{u}_k^*\|_2 \leq \epsilon;\ \forall k \in [K]$$

[a] This is a standard assumption to guarantee the convergence of the power iteration, as in Golub and Van Loan [1996]

Reference

Gene H. Golub and Charles F. Van Loan. *Matrix Computations*. The Johns Hopkins University Press, third edition, 1996.

Idea of Proof

Lemma 1 (Davis-Kahan $\sin\Theta$ Theorem). *Let $\mathbf{M}^* \in \mathbb{R}^{d\times d}$ and let $\mathbf{M} = \mathbf{M}^* + \mathbf{H}$. Let $\mathbf{a}_1^*$ and $\mathbf{a}_1$ be the top eigenvectors of $\mathbf{M}^*$ and $\mathbf{M}$, respectively. Then we have:*

$$\sin\angle\{\mathbf{a}_1^*, \mathbf{a}_1\} \leq \frac{\|\mathbf{H}\|_2}{\min_{j\neq k}|\sigma_k^* - \sigma_j|}.$$

Based on Lemma 1, we show that

- $\|\mathbf{u}_k - \mathbf{u}_k^*\|_2 \leq \frac{2}{\mathcal{T}_k}\|\mathbf{\Sigma}_k - \mathbf{\Sigma}_k^*\|_F$
- $\|\mathbf{\Sigma}_{k+1} - \mathbf{\Sigma}_{k+1}^*\|_F \leq \left(3 + \frac{2}{\mathcal{T}_k}\right)\|\mathbf{u}_k - \mathbf{u}_k^*\|_2 + 2\|\boldsymbol{\delta}_k\|_2$

Analysis of Power Iteration

Building on top of previous analysis, we note that

- $\left\|\left(\mathbf{\Sigma}_k - \mathbf{\Sigma}_k^*\right)\mathbf{u}_j^*\right\|_2$ evolves differently for each j.
- $\left|\left(\mathbf{u}_k - \mathbf{u}_k^*\right)^\top \mathbf{u}_j^*\right|$ evolves differently for each j.

Lemma 2. *Let $\mathbf{v}_k$ be the output of the k-th power iteration. For all $k \in [d]$ and $j \geq k$, we have:*

$$\left\|\left(\mathbf{\Sigma}_k^* - \mathbf{\Sigma}_k\right)\mathbf{u}_j^*\right\|_2 \leq \sum_{k'=1}^{k-1} \lambda_{k'}\left|\mathbf{v}_{k'}^\top \mathbf{u}_j^*\right|. \tag{8}$$

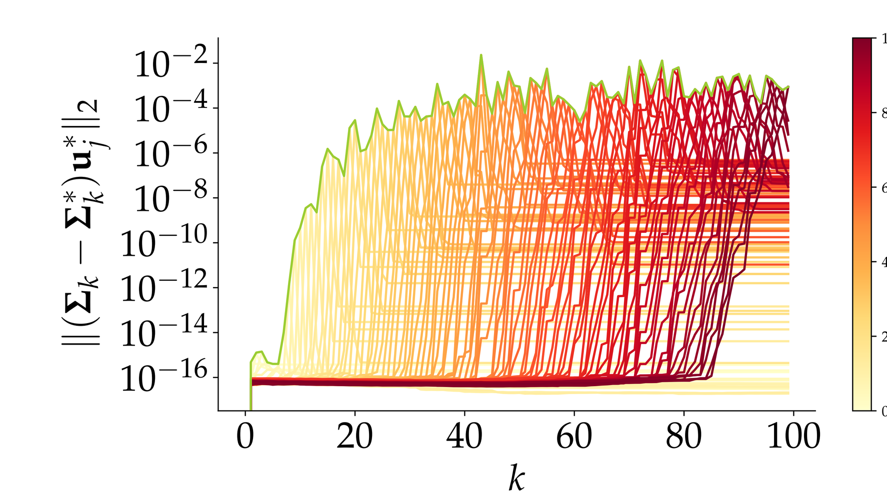

Lemma 3. *Let $\mathbf{x}_t$ be the result of running power iteration starting from $\mathbf{x}_0$ for t iterations, as defined in (5). Let $\mathbf{M} = \mathbf{M}^* + \mathbf{H}$. Moreover, let $\sigma_j^*, \mathbf{a}_j^*$ be the j-th eigenvalue and eigenvector of $\mathbf{M}^*$, and let σ_j be the j-th eigenvalue of $\mathbf{M}$. Assume that there exists $c_0 \in (0, \infty)$ such that $\left|\mathbf{x}_0^\top \mathbf{a}_1\right| \geq c_0^{-1}$. Then, for all $j = 2, \ldots, r$, we have:*

$$\left|\mathbf{x}_t^\top \mathbf{a}_j^*\right| \leq c_0\left(\left(\frac{\sigma_j^*}{\sigma_1}\right)^t + \frac{\sigma_j^*}{\sigma_1 - \sigma_j^*}\|\mathbf{H}\mathbf{a}_j^*\|_2\right).$$