



Mathematisches
Forschungsinstitut
Oberwolfach



Oberwolfach Preprints

OWP 2018 - 12

ARMIN EFTEKHARI, RAPHAEL A. HAUSER AND
ANDREAS GRAMMENOS

MOSES: A Streaming Algorithm for Linear
Dimensionality Reduction

Mathematisches Forschungsinstitut Oberwolfach gGmbH
Oberwolfach Preprints (OWP) ISSN 1864-7596

Oberwolfach Preprints (OWP)

Starting in 2007, the MFO publishes a preprint series which mainly contains research results related to a longer stay in Oberwolfach. In particular, this concerns the Research in Pairs-Programme (RiP) and the Oberwolfach-Leibniz-Fellows (OWLF), but this can also include an Oberwolfach Lecture, for example.

A preprint can have a size from 1 - 200 pages, and the MFO will publish it on its website as well as by hard copy. Every RiP group or Oberwolfach-Leibniz-Fellow may receive on request 30 free hard copies (DIN A4, black and white copy) by surface mail.

Of course, the full copy right is left to the authors. The MFO only needs the right to publish it on its website *www.mfo.de* as a documentation of the research work done at the MFO, which you are accepting by sending us your file.

In case of interest, please send a **pdf file** of your preprint by email to *rip@mfo.de* or *owlf@mfo.de*, respectively. The file should be sent to the MFO within 12 months after your stay as RiP or OWLF at the MFO.

There are no requirements for the format of the preprint, except that the introduction should contain a short appreciation and that the paper size (respectively format) should be DIN A4, "letter" or "article".

On the front page of the hard copies, which contains the logo of the MFO, title and authors, we shall add a running number (20XX - XX).

We cordially invite the researchers within the RiP or OWLF programme to make use of this offer and would like to thank you in advance for your cooperation.

Imprint:

Mathematisches Forschungsinstitut Oberwolfach gGmbH (MFO)
Schwarzwaldstrasse 9-11
77709 Oberwolfach-Walke
Germany

Tel +49 7834 979 50
Fax +49 7834 979 55
Email admin@mfo.de
URL www.mfo.de

The Oberwolfach Preprints (OWP, ISSN 1864-7596) are published by the MFO.
Copyright of the content is held by the authors.

DOI 10.14760/OWP-2018-12

MOSES: A Streaming Algorithm for Linear Dimensionality Reduction

Armin Eftekhari, Raphael A. Hauser, Andreas Grammenos*

June 12, 2018

Abstract

This paper introduces Memory-limited Online Subspace Estimation Scheme (MOSES) for both estimating the principal components of data and reducing its dimension. More specifically, consider a scenario where the data vectors are presented sequentially to a user who has limited storage and processing time available, for example in the context of sensor networks. In this scenario, MOSES maintains an estimate of leading principal components of the data that has arrived so far and also reduces its dimension. In terms of its origins, MOSES slightly generalises the popular incremental Singular Value Decomposition (SVD) to handle thin blocks of data. This simple generalisation is in part what allows us to complement MOSES with a comprehensive statistical analysis that is not available for incremental SVD, despite its empirical success. This generalisation also enables us to concretely interpret MOSES as an approximate solver for the underlying non-convex optimisation program. We also find that MOSES shows state-of-the-art performance in our numerical experiments with both synthetic and real-world datasets.

Keywords— Principal component analysis, Linear dimensionality reduction, Subspace identification, Streaming algorithms, Non-convex optimisation.

1 Introduction

Linear models are pervasive in data and computational sciences and Principal Component Analysis (PCA) in particular is an indispensable tool for detecting linear structure in collected data [1, 2, 3, 4, 5, 6, 7]. Principal components are the directions that preserve most of the “energy” of a dataset and can be used for linear dimensionality reduction, among other things. In turn, successful dimensionality reduction is at the heart of classification, regression, and other learning tasks that often suffer from the “curse of dimensionality”, where having a small number of training samples in relation to the dimension of data typically leads to overfitting [8].

In this work, we are interested in both computing the principal components *and* reducing the dimension of data that is presented sequentially to a user. Due to hardware limitations, the user can only store small amounts of data, which in turn severely limits the available processing time for each incoming data vector. For example, consider a network of battery-powered and cheap sensors that must relay their measurements to a central node on a daily basis. Each sensor has a small storage and does not have the power to relay all the raw data to the central node. One solution is then for each sensor to reduce the dimension of its data to make transmission to the central node possible. Even if each sensor had unlimited storage, the frequent daily updates scheduled by the central node would force each sensor to reduce the dimension of its data “on the go” before transmitting it to the central node. A number of similar problems are listed in [9].

Motivated by such scenarios, we are interested in developing a *streaming* algorithm for linear dimensionality reduction, namely an algorithm with minimal storage and computational requirements. As more and more data vectors arrive, this algorithm would keep a running estimate of the principal components of the

*AE is with the Alan Turing Institute in London and the School of Mathematics, University of Edinburgh. RAH is with the Mathematical Institute at the University of Oxford and the Alan Turing Institute. AG is with the University of Cambridge and the Alan Turing Institute. Emails: aeftekhari@turing.ac.uk, hauser@maths.ox.ac.uk, ag926@c1.cam.ac.uk.

data *and* project the available data onto this estimate to reduce its dimension. As we will see shortly, what we need here is a streaming algorithm to compute truncated Singular Value Decomposition (SVD).

Incremental SVD is a streaming algorithm that updates its estimate of truncated SVD of the data matrix with every new incoming vector [10, 11, 12, 13, 14]. To the best of our knowledge and despite its popularity, incremental SVD lacks comprehensive statistical guarantees. In fact, [15] only very recently provided stochastic analysis for two of the variants of incremental SVD in [16, 17]. To be specific, in [15] the authors studied how well the output of incremental SVD approximates (only) the leading principal component of data, in expectation. In particular, [15] does *not* offer any guarantees for dimensionality reduction, see Section 5 for a detailed review of the prior art.

Contributions. In this paper, we propose Memory-limited Online Subspace Estimation Scheme (MOSES) for streaming dimensionality reduction. MOSES slightly generalises incremental SVD to update its estimate with every incoming *thin* block of data, rather than with every incoming vector. This small difference between incremental SVD and MOSES is in part what enables us to complement MOSES with a comprehensive statistical analysis. Indeed, Theorem 1 below considers the important case where the incoming data vectors are drawn from a zero-mean normal distribution. This stochastic setup is a powerful generalisation of the popular *spiked covariance* model common in statistical signal processing [18]. Theorem 1 states that MOSES nearly matches the performance of “offline” truncated SVD (which has unlimited memory and computing resources), provided that the corresponding covariance matrix is well-conditioned and has a small residual. Moreover, we concretely interpret MOSES as an approximate solver for the underlying non-convex optimisation program. We also find that MOSES shows state-of-the-art performance in our numerical experiments with both synthetic and real-world datasets.

Organisation. The rest of this paper is organised as follows. Section 2 formally introduces MOSES and Section 3 motivates this algorithm from an optimisation viewpoint. Section 4 then presents the statistical guarantees for MOSES, summarised in Theorem 1. Prior art is reviewed in Section 5 and Section 6 numerically compares MOSES with two main alternatives in the current literature. The proofs are deferred to the supplementary material.

2 Introducing MOSES

Consider a sequence of vectors $\{y_t\}_{t=1}^T \subset \mathbb{R}^n$, presented to us sequentially, and let

$$\mathbf{Y}_T := [y_1 \quad y_2 \quad \cdots \quad y_T] \in \mathbb{R}^{n \times T}, \quad (1)$$

for short. We conveniently assume throughout that \mathbf{Y}_T is *centred*, namely the entries of each row of \mathbf{Y}_T sum up to zero. It is then a consequence of the celebrated Eckart-Young-Mirsky Theorem [19, 20] that leading r principal components of \mathbf{Y}_T in fact coincide with leading r left singular vectors of \mathbf{Y}_T . More specifically, let

$$\mathbf{Y}_T \stackrel{\text{SVD}}{=} \mathbf{S}_T \mathbf{\Gamma}_T \mathbf{Q}_T^* \quad (2)$$

be the SVD of \mathbf{Y}_T , where $\mathbf{S}_T \in \mathbb{R}^{n \times n}$ and $\mathbf{Q}_T \in \mathbb{R}^{T \times T}$ are orthonormal bases, and the diagonal matrix $\mathbf{\Gamma}_T \in \mathbb{R}^{n \times T}$ contains the singular values of \mathbf{Y}_T in nonincreasing order. Let us assume that $\text{rank}(\mathbf{Y}_T) \geq r$. Then the first r columns of \mathbf{S}_T , which we collect in $\mathbf{S}_{T,r} \in \mathbb{R}^{n \times r}$, are leading r principal components of \mathbf{Y}_T . We accordingly decompose \mathbf{Y}_T into two components, namely

$$\begin{aligned} \mathbf{Y}_T &\stackrel{\text{SVD}}{=} \mathbf{S}_T \mathbf{\Gamma}_T \mathbf{Q}_T^* \\ &= [\mathbf{S}_{T,r} \quad \mathbf{S}_{T,r^+}] \begin{bmatrix} \mathbf{\Gamma}_{T,r} & \\ & \mathbf{\Gamma}_{T,r^+} \end{bmatrix} [\mathbf{Q}_{T,r} \quad \mathbf{Q}_{T,r^+}^*] \\ &= \mathbf{S}_{T,r} \mathbf{\Gamma}_{T,r}^* \mathbf{Q}_{T,r} + \mathbf{S}_{T,r^+} \mathbf{\Gamma}_{T,r^+} \mathbf{Q}_{T,r^+}^* \\ &=: \mathbf{Y}_{T,r} + \mathbf{Y}_{T,r^+}, \end{aligned} \quad (3)$$

where the empty blocks are zero. It is easy to see that the column and row spaces of $\mathbf{Y}_{T,r}$ are orthogonal to those of $\mathbf{Y}_{T,r+}$, namely

$$\mathbf{Y}_{T,r} \mathbf{Y}_{T,r+}^* = \mathbf{0}_{n \times n}, \quad \mathbf{Y}_{T,r}^* \mathbf{Y}_{T,r+} = \mathbf{0}_{T \times T}. \quad (4)$$

Moreover, the Eckart-Young-Mirsky Theorem implies that $\mathbf{Y}_{T,r} = \text{SVD}_r(\mathbf{Y}_T)$ is a rank- r truncation of \mathbf{Y}_T . That is, $\mathbf{Y}_{T,r}$ is a best rank- r approximation of \mathbf{Y}_T with the corresponding residual

$$\begin{aligned} \|\mathbf{Y}_T - \mathbf{Y}_{T,r}\|_F^2 &= \min_{\text{rank}(\mathbf{X})=r} \|\mathbf{Y}_T - \mathbf{X}\|_F^2 \\ &= \|\mathbf{Y}_{T,r+}\|_F^2 \quad (\text{see (4)}) \\ &= \sum_{i \geq r+1} \sigma_i^2(\mathbf{Y}_T) \\ &=: \rho_r^2(\mathbf{Y}_T), \end{aligned} \quad (5)$$

where $\sigma_1(\mathbf{Y}_T) \geq \sigma_2(\mathbf{Y}_T) \geq \dots$ are the singular values of \mathbf{Y}_T . Given leading r principal components of \mathbf{Y}_T , namely $\mathbf{S}_{T,r}$ in (3), we can reduce the dimension of data from n to r by projecting \mathbf{Y}_T onto the span of $\mathbf{S}_{T,r}$, that is

$$\begin{aligned} \mathbf{S}_{T,r}^* \cdot \mathbf{Y}_T &= \mathbf{S}_{T,r}^* (\mathbf{Y}_{T,r} + \mathbf{Y}_{T,r+}) \quad (\text{see (3)}) \\ &= \mathbf{S}_{T,r}^* \mathbf{Y}_{T,r} \quad (\text{see (4)}) \\ &= \mathbf{\Gamma}_{T,r} \mathbf{Q}_{T,r}^* \in \mathbb{R}^{r \times T}. \quad (\text{see (4)}) \end{aligned} \quad (6)$$

The projected data matrix $\mathbf{S}_{T,r}^* \mathbf{Y}_T \in \mathbb{R}^{r \times T}$ again has T data vectors (namely, columns) but these vectors are embedded in (often much smaller) \mathbb{R}^r rather than \mathbb{R}^n . Note also that

$$\begin{aligned} \mathbf{Y}_{T,r} &= \text{SVD}_r(\mathbf{Y}_T) \\ &= \mathbf{S}_{T,r} \mathbf{S}_{T,r}^* \cdot \mathbf{Y}_T \\ &= \mathbf{S}_{T,r} \mathbf{S}_{T,r}^* (\mathbf{Y}_{T,r} + \mathbf{Y}_{T,r+}) \quad (\text{see (3)}) \\ &= \mathbf{S}_{T,r} \mathbf{S}_{T,r}^* \mathbf{Y}_{T,r} \quad (\text{see (4)}) \\ &= \underbrace{\mathbf{S}_{T,r}}_{\text{PCs}} \cdot \underbrace{\mathbf{\Gamma}_{T,r} \mathbf{Q}_{T,r}^*}_{\text{projected data}}. \quad (\text{see (3)}) \end{aligned} \quad (7)$$

That is, rank- r truncation of \mathbf{Y}_T encapsulates both leading r principal components of \mathbf{Y}_T , namely $\mathbf{S}_{T,r}$, and the projected data matrix $\mathbf{S}_{T,r}^* \mathbf{Y}_T = \mathbf{\Gamma}_{T,r} \mathbf{Q}_{T,r}^*$. In other words, computing a rank- r truncation of the data matrix both yields its principal components and reduces the dimension of data at once.

We are in this work interested in developing a streaming algorithm to compute $\mathbf{Y}_{T,r} = \text{SVD}_r(\mathbf{Y}_T)$, a rank- r truncation of the data matrix \mathbf{Y}_T . More specifically, to compute $\mathbf{Y}_{T,r}$, we are only allowed one pass through the columns of \mathbf{Y}_T and have available a limited amount of storage, namely $O(n)$ bits. See also Figure 1.

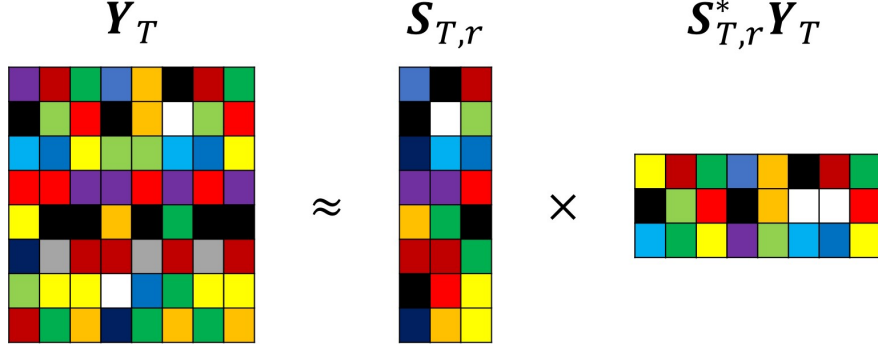


Figure 1: Given a data matrix $\mathbf{Y}_T \in \mathbb{R}^{n \times T}$, truncated SVD finds the best low-dimensional linear model to represent the data: For a typically small integer r , we compute $\mathbf{Y}_{T,r} = \text{SVD}_r(\mathbf{Y}_T) = \mathbf{S}_{T,r} \cdot \mathbf{S}_{T,r}^* \mathbf{Y}_T$, where $\mathbf{S}_{T,r} \in \mathbb{R}^{n \times r}$ contains leading r principal components of \mathbf{Y}_T and $\mathbf{S}_{T,r}^* \mathbf{Y}_T \in \mathbb{R}^{r \times T}$ is the projected data matrix with reduced dimension r (instead of n). This paper presents MOSES, a streaming algorithm for truncated SVD. Put differently, MOSES keeps both a running estimate of the principal components and the projection of data, received so far, onto this estimate.

For a block size b , our strategy is to iteratively group every b incoming vectors into an $n \times b$ block and then update a rank- r estimate of the data that has been received so far. We assume throughout that $r \leq b \leq T$ and in fact often take the block size as $b = O(r)$. It is convenient to assume that the number of blocks $K := T/b$ is an integer. Upon arrival of a new data block $\{y_t\}_{t=(k-1)b+1}^{kb}$, we concatenate these vectors to form the matrix

$$\mathbf{y}_k = \begin{bmatrix} y_{(k-1)b+1} & \cdots & y_{kb} \end{bmatrix} \in \mathbb{R}^{n \times b}.$$

For every $k \in [1 : K] := \{1, \dots, K\}$, we then set

$$\hat{\mathbf{Y}}_{kb,r} = \text{SVD}_r \left(\begin{bmatrix} \hat{\mathbf{Y}}_{(k-1)b,r} & \mathbf{y}_k \end{bmatrix} \right) \in \mathbb{R}^{n \times kb}, \quad (8)$$

with the convention that $\hat{\mathbf{Y}}_{0,r}$ is the empty matrix. We call this simple algorithm MOSES for Memory-limited Online Subspace Estimation Scheme. The output of MOSES after K iterations is

$$\hat{\mathbf{Y}}_{Kb,r} = \hat{\mathbf{Y}}_{T,r},$$

which contains both an estimate of leading r principal components of \mathbf{Y}_T and the projection of \mathbf{Y}_T onto this estimate, as discussed below. For easy reference, MOSES is summarised in Algorithm 1. An efficient implementation of MOSES is given in Algorithm 2, which explicitly maintains both the estimates of principal components and the projected data. As discussed below, the storage and computational requirements of Algorithm 2 are nearly optimal.

Discussion. MOSES maintains a rank- r estimate of the data that has been received so far, and updates its estimate in every iteration to account for the new incoming block of data. In other words, MOSES simultaneously keeps an estimate of principal components *and* the projection of the available data onto this estimate. More specifically, note that the final output of MOSES, namely $\hat{\mathbf{Y}}_{T,r} \in \mathbb{R}^{n \times T}$, is at most rank- r , and let

$$\hat{\mathbf{Y}}_{T,r} \stackrel{\text{SVD}}{=} \hat{\mathbf{S}}_{T,r} \hat{\mathbf{\Gamma}}_{T,r} \hat{\mathbf{Q}}_{T,r}^*, \quad (9)$$

be its SVD. Then, $\hat{\mathbf{S}}_{T,r} \in \mathbb{R}^{n \times r}$ is MOSES's estimate of principal components of the data matrix \mathbf{Y}_T , and

$$\hat{\mathbf{S}}_{T,r}^* \hat{\mathbf{Y}}_{T,r} = \hat{\mathbf{\Gamma}}_{T,r} \hat{\mathbf{Q}}_{T,r} \in \mathbb{R}^{r \times T}$$

is the projection of $\hat{\mathbf{Y}}_{T,r}$ onto this estimate. That is, $\hat{\mathbf{S}}_{T,r}^* \hat{\mathbf{Y}}_{T,r}$ is the MOSES's estimate of the projected data matrix.

It is natural to ask how MOSES compares with the “offline” truncated SVD. To be concrete, recall that $\mathbf{Y}_{T,r} = \text{SVD}_r(\mathbf{Y}_T)$ is a rank- r truncation of the data matrix \mathbf{Y}_T with the corresponding residual of $\rho_r^2(\mathbf{Y}_T)$, see (5). Because $\mathbf{Y}_{T,r}$ is a best rank- r approximation of \mathbf{Y}_T , the final output $\widehat{\mathbf{Y}}_{T,r}$ of MOSES cannot be a better approximation of \mathbf{Y}_T , that is

$$\min_{\text{rank}(\mathbf{X})=r} \|\mathbf{Y}_T - \mathbf{X}\|_F^2 = \|\mathbf{Y}_T - \mathbf{Y}_{T,r}\|_F^2 = \rho_r^2(\mathbf{Y}_T) \leq \|\mathbf{Y}_T - \widehat{\mathbf{Y}}_{T,r}\|_F^2. \quad (10)$$

However, our main technical contribution in Theorem 1 below states that, under certain conditions, $\widehat{\mathbf{Y}}_{T,r}$ is not much worse than $\mathbf{Y}_{T,r}$, in the sense that

$$\rho_r^2(\mathbf{Y}_T) \leq \|\mathbf{Y}_T - \widehat{\mathbf{Y}}_{T,r}\|_F^2 \lesssim \text{poly}(T) \cdot \rho_r^2(\mathbf{Y}_T), \quad (11)$$

and the polynomial factor above is relatively small. That is, MOSES for streaming dimensionality reduction nearly matches the performance of its offline version that has access to unlimited storage and computing resources, see Section 4 for the details.

Origins. Incremental SVD is a streaming algorithm that updates its estimate of (truncated) SVD of the data matrix with every new incoming vector [10, 11, 12, 13, 14]. It is easy to verify that MOSES slightly generalises incremental SVD to update its estimate with every incoming block of data, rather than with every incoming data vector. As detailed later in Section 5, this small difference between incremental SVD and MOSES is in part what enables us to complement MOSES with a comprehensive statistical analysis in Theorem 1 which is, to the best of our knowledge, absent from the literature of incremental SVD, despite its popularity and empirical success. See Section 5 for more details. This small difference also allows us to concretely interpret MOSES as an approximate solver for the underlying non-convex program, as detailed in Section 3.

Storage and computational requirements. The efficient implementation of MOSES in Algorithm 2 is based on the ideas from incremental SVD and it is straightforward to verify that Algorithms 1 and 2 are indeed equivalent; at iteration k , the relation between the output of Algorithm 1 ($\widehat{\mathbf{Y}}_{kb,r}$) and the output of Algorithm 2 ($\widehat{\mathbf{S}}_{kb,r}, \widehat{\mathbf{\Gamma}}_{kb,r}, \widehat{\mathbf{Q}}_{kb,r}$) is

$$\widehat{\mathbf{Y}}_{kb,r} \stackrel{\text{SVD}}{=} \widehat{\mathbf{S}}_{kb,r} \widehat{\mathbf{\Gamma}}_{kb,r} \widehat{\mathbf{Q}}_{kb,r}^*,$$

where the right-hand side above is the SVD of $\widehat{\mathbf{Y}}_{kb,r}$. More specifically, $\widehat{\mathbf{S}}_{kb,r} \in \mathbb{R}^{n \times r}$ has orthonormal columns and is the MOSES’s estimate of leading r principal components of $\mathbf{Y}_{kb} \in \mathbb{R}^{n \times kb}$, where we recall that \mathbf{Y}_{kb} is the data received so far. Moreover,

$$\widehat{\mathbf{S}}_{kb,r}^* \widehat{\mathbf{Y}}_{kb,r} = \widehat{\mathbf{\Gamma}}_{kb,r} \widehat{\mathbf{Q}}_{kb,r}^* \in \mathbb{R}^{r \times kb}$$

is the projection of $\widehat{\mathbf{Y}}_{kb,r}$ onto this estimate, namely $\widehat{\mathbf{S}}_{kb,r}^* \widehat{\mathbf{Y}}_{kb,r}$ is MOSES’s estimate of the projected data matrix so far. In words, the efficient implementation of MOSES in Algorithm 2 explicitly maintains estimates of both principal components and the projected data, at every iteration.

Let us now evaluate the storage and computational requirements of MOSES. At the start of iteration k , Algorithm 2 stores the matrices

$$\widehat{\mathbf{S}}_{(k-1)b,r} \in \mathbb{R}^{n \times r}, \quad \widehat{\mathbf{\Gamma}}_{(k-1)b,r} \in \mathbb{R}^{r \times r}, \quad \widehat{\mathbf{Q}}_{(k-1)b,r} \in \mathbb{R}^{(k-1)b \times r},$$

and after that also receives and stores the incoming block $\mathbf{y}_k \in \mathbb{R}^{n \times b}$. This requires $O(r(n + (k-1)b + 1)) + O(bn)$ bits of memory, because $\widehat{\mathbf{\Gamma}}_{(k-1)b,r}$ is diagonal. Assuming that $b = O(r)$, Algorithm 2 therefore requires $O(r(n + kr))$ bits of memory at iteration k . Note that this is optimal, as it is impossible to store a rank- r matrix of size $n \times kb$ with fewer bits when $b = O(r)$.

It is also easy to verify that Algorithm 2 performs $O(r^2(n + kb)) = O(r^2(n + kr))$ flops in iteration k . The dependence of both storage and computational complexity on k is due to the fact that MOSES maintains both an estimate of principal components in $\widehat{\mathbf{S}}_{kb,r}$ and an estimate of the projected data in $\widehat{\mathbf{\Gamma}}_{kb,r} \widehat{\mathbf{Q}}_{kb,r}^*$. To

maximise the efficiency, one might optionally “flush out” the projected data after every n/b iterations, as described in the last step in Algorithm 2.

Algorithm 1 MOSES: A streaming algorithm for linear dimensionality reduction

Input: Sequence of vectors $\{y_t\}_{t \geq 1} \subset \mathbb{R}^n$, rank r , and block size $b \geq r$.

Output: Sequence $\{\widehat{\mathbf{Y}}_{kb,r}\}_k$, where $\widehat{\mathbf{Y}}_{kb,r} \in \mathbb{R}^{n \times kb}$ for every $k \geq 1$.

Body:

1. Set $\widehat{\mathbf{Y}}_{0,r} \leftarrow \{\}$.
 2. For $k \geq 1$, repeat
 - (a) Form $\mathbf{y}_k \in \mathbb{R}^{n \times b}$ by concatenating $\{y_t\}_{t=(k-1)b+1}^{kb}$.
 - (b) Set $\widehat{\mathbf{Y}}_{kb,r} = \text{SVD}_r([\widehat{\mathbf{Y}}_{(k-1)b,r} \ \mathbf{y}_k])$, where $\text{SVD}_r(\cdot)$ returns a rank- r truncated SVD of its argument.
-

Algorithm 2 Efficient implementation of MOSES in Algorithm 1

Input: Sequence of vectors $\{y_t\}_t \subset \mathbb{R}^n$ and block size b .

Output: Sequence $\{\widehat{\mathbf{S}}_{kb,r}, \widehat{\mathbf{\Gamma}}_{kb,r}, \widehat{\mathbf{Q}}_{kb,r}\}_k$.

Body:

1. For $k = 1$,

(a) Form $\mathbf{y}_1 \in \mathbb{R}^{n \times b}$ by concatenating $\{y_t\}_{t=1}^b$.

(b) Set

$$[\widehat{\mathbf{S}}_{b,r}, \widehat{\mathbf{\Gamma}}_{b,r}, \widehat{\mathbf{Q}}_{b,r}] = \text{SVD}_r(\mathbf{y}_1),$$

where $\widehat{\mathbf{S}}_{b,r} \in \mathbb{R}^{n \times r}$ and $\widehat{\mathbf{Q}}_{b,r} \in \mathbb{R}^{b \times r}$ have orthonormal columns, and the diagonal matrix $\widehat{\mathbf{\Gamma}}_{b,r} \in \mathbb{R}^{r \times r}$ contains leading r singular values.

2. For $k \geq 2$, repeat

(a) Form $\mathbf{y}_k \in \mathbb{R}^{n \times b}$ by concatenating $\{y_t\}_{t=(k-1)b+1}^{kb}$.

(b) Set

$$\begin{aligned} \dot{\mathbf{q}}_k &= \widehat{\mathbf{S}}_{(k-1)b,r}^* \mathbf{y}_k \in \mathbb{R}^{r \times b}, \\ \widehat{\mathbf{z}}_k &= \mathbf{y}_k - \widehat{\mathbf{S}}_{(k-1)b,r} \dot{\mathbf{q}}_k \in \mathbb{R}^{n \times b}. \end{aligned}$$

(c) Let $[\widehat{\mathbf{s}}_k, \mathbf{v}_k] = \text{QR}(\widehat{\mathbf{z}}_k)$ be the QR decomposition of $\widehat{\mathbf{z}}_k$, where $\widehat{\mathbf{s}}_k \in \mathbb{R}^{n \times b}$ has orthonormal columns and $\mathbf{v}_k \in \mathbb{R}^{b \times b}$.

(d) Let

$$[\mathbf{u}_k, \widehat{\mathbf{\Gamma}}_{kb,r}, \widehat{\mathbf{q}}_k] = \text{SVD}_r \left(\begin{bmatrix} \widehat{\mathbf{\Gamma}}_{(k-1)b,r} & \dot{\mathbf{q}}_k \\ \mathbf{0}_{b \times r} & \mathbf{v}_k \end{bmatrix} \right), \quad (12)$$

where $\mathbf{u}_k, \widehat{\mathbf{q}}_k \in \mathbb{R}^{(r+b) \times r}$ have orthonormal columns and the diagonal matrix $\widehat{\mathbf{\Gamma}}_{kb,r} \in \mathbb{R}^{r \times r}$ contains leading r singular values in nonincreasing order.

(e) Let

$$\widehat{\mathbf{S}}_{kb,r} = \begin{bmatrix} \widehat{\mathbf{S}}_{(k-1)b,r} & \widehat{\mathbf{s}}_k \end{bmatrix} \mathbf{u}_k.$$

(f) **(optional)** If the number of rows of $\widehat{\mathbf{Q}}_{(k-1)b,r}$ exceeds n and $\widehat{\mathbf{Q}}_{(k-1)b,r}$ is not needed any more, it is optional in order to improve efficiency to set $\widehat{\mathbf{Q}}_{kb,r} = \widehat{\mathbf{q}}_k$.

(g) Otherwise, set

$$\widehat{\mathbf{Q}}_{kb,r} = \begin{bmatrix} \widehat{\mathbf{Q}}_{(k-1)b,r} & \mathbf{0} \\ \mathbf{0} & \mathbf{I}_b \end{bmatrix} \widehat{\mathbf{q}}_k. \quad (13)$$

3 Optimisation Viewpoint

MOSES has a natural interpretation as an approximate solver for the non-convex optimisation program underlying PCA, which serves as its motivation. More specifically, recall that leading r principal components of \mathbf{Y}_T are obtained by solving the non-convex program

$$\min_{\mathcal{U} \in \mathbb{G}(n,r)} \|\mathbf{Y}_T - \mathbf{P}_{\mathcal{U}} \mathbf{Y}_T\|_F^2, \quad (14)$$

where the minimization is over the Grassmannian $\mathbb{G}(n,r)$, the set of all r -dimensional subspaces in \mathbb{R}^n . Above, $\mathbf{P}_{\mathcal{U}} \in \mathbb{R}^{n \times n}$ is the orthogonal projection onto the subspace \mathcal{U} . By construction in Section 2, note

that

$$\begin{aligned} \mathbf{Y}_T &= [y_1 \quad y_2 \quad \cdots \quad y_T] \quad (\text{see (1)}) \\ &= [\mathbf{y}_1 \quad \mathbf{y}_2 \quad \cdots \quad \mathbf{y}_K] \in \mathbb{R}^{n \times T}, \end{aligned} \quad (15)$$

where $\{\mathbf{y}_k\}_{k=1}^K$ are the incoming blocks of data. This allows us to rewrite Program (14) as

$$\begin{aligned} \min_{\mathcal{U} \in \mathbb{G}(n,r)} \|\mathbf{Y}_T - \mathbf{P}_{\mathcal{U}} \mathbf{Y}_T\|_F^2 &= \min_{\mathcal{U} \in \mathbb{G}(n,r)} \sum_{k=1}^K \|\mathbf{y}_k - \mathbf{P}_{\mathcal{U}} \mathbf{y}_k\|_F^2 \quad (\text{see (15)}) \\ &= \begin{cases} \min \sum_{k=1}^K \|\mathbf{y}_k - \mathbf{P}_{\mathcal{U}_K} \cdots \mathbf{P}_{\mathcal{U}_1} \mathbf{y}_k\|_F^2 \\ \mathcal{U}_1 = \mathcal{U}_2 = \cdots = \mathcal{U}_K, \end{cases} \end{aligned} \quad (16)$$

where the last minimisation above is over all identical subspaces $\{\mathcal{U}_k\}_{k=1}^K \subset \mathbb{G}(n,r)$. Our strategy is to make a sequence of approximations to the program in the last line above. In the first approximation, we only keep the first summand in the last line of (16). That is, our first approximation reads as

$$\begin{aligned} \begin{cases} \min \sum_{k=1}^K \|\mathbf{y}_k - \mathbf{P}_{\mathcal{U}_K} \cdots \mathbf{P}_{\mathcal{U}_1} \mathbf{y}_k\|_F^2 \\ \mathcal{U}_1 = \mathcal{U}_2 = \cdots = \mathcal{U}_K \end{cases} &\geq \begin{cases} \min \|\mathbf{y}_1 - \mathbf{P}_{\mathcal{U}_K} \cdots \mathbf{P}_{\mathcal{U}_1} \mathbf{y}_1\|_F^2 \\ \mathcal{U}_1 = \mathcal{U}_2 = \cdots = \mathcal{U}_K \end{cases} \\ &= \min_{\mathcal{U} \in \mathbb{G}(n,r)} \|\mathbf{y}_1 - \mathbf{P}_{\mathcal{U}} \mathbf{y}_1\|_F^2, \end{aligned} \quad (17)$$

where the second line above follows by setting $\mathcal{U} = \mathcal{U}_1 = \cdots = \mathcal{U}_K$. Let $\widehat{\mathcal{S}}_{b,r}$ be a minimiser of the program in the last line above. Note that $\widehat{\mathcal{S}}_{b,r}$ simply spans leading r principal components of \mathbf{y}_1 , akin to Program (14). This indeed coincides with the output of MOSES in the first iteration, because

$$\begin{aligned} \widehat{\mathbf{Y}}_{b,r} &= \text{SVD}_r(\mathbf{y}_1) \quad (\text{see (8)}) \\ &= \mathbf{P}_{\widehat{\mathcal{S}}_{b,r}} \mathbf{y}_1. \quad (\text{similar to the second line of (7)}) \end{aligned} \quad (18)$$

Next consider the next approximation in which we keep two of the summands in the last line of (16), namely

$$\begin{cases} \min \sum_{k=1}^K \|\mathbf{y}_k - \mathbf{P}_{\mathcal{U}_K} \cdots \mathbf{P}_{\mathcal{U}_1} \mathbf{y}_k\|_F^2 \\ \mathcal{U}_1 = \mathcal{U}_2 = \cdots = \mathcal{U}_K \end{cases} \geq \begin{cases} \min \|\mathbf{y}_1 - \mathbf{P}_{\mathcal{U}_K} \cdots \mathbf{P}_{\mathcal{U}_1} \mathbf{y}_1\|_F^2 + \|\mathbf{y}_2 - \mathbf{P}_{\mathcal{U}_K} \cdots \mathbf{P}_{\mathcal{U}_1} \mathbf{y}_2\|_F^2 \\ \mathcal{U}_1 = \mathcal{U}_2 = \cdots = \mathcal{U}_K, \end{cases} \quad (19)$$

and then we substitute $\mathcal{U}_1 = \widehat{\mathcal{S}}_{b,r}$ above to arrive at the new program

$$\begin{aligned} &\begin{cases} \min \|\mathbf{y}_1 - \mathbf{P}_{\mathcal{U}_K} \cdots \mathbf{P}_{\mathcal{U}_2} \mathbf{P}_{\widehat{\mathcal{S}}_{b,r}} \mathbf{y}_1\|_F^2 + \|\mathbf{y}_2 - \mathbf{P}_{\mathcal{U}_K} \cdots \mathbf{P}_{\mathcal{U}_2} \mathbf{y}_2\|_F^2 \\ \mathcal{U}_2 = \mathcal{U}_3 = \cdots = \mathcal{U}_K \end{cases} \\ &= \min_{\mathcal{U} \in \mathbb{G}(n,r)} \|\mathbf{y}_1 - \mathbf{P}_{\mathcal{U}} \mathbf{P}_{\widehat{\mathcal{S}}_{b,r}} \mathbf{y}_1\|_F^2 + \|\mathbf{y}_2 - \mathbf{P}_{\mathcal{U}} \mathbf{y}_2\|_F^2, \end{aligned} \quad (20)$$

where the second program above follows by setting $\mathcal{U} = \mathcal{U}_2 = \dots = \mathcal{U}_K$. We can rewrite the above program as

$$\begin{aligned}
& \min_{\mathcal{U} \in \mathbb{G}(n,r)} \|\mathbf{y}_1 - \mathbf{P}_{\mathcal{U}} \mathbf{P}_{\widehat{\mathcal{S}}_{b,r}} \mathbf{y}_1\|_F^2 + \|\mathbf{y}_2 - \mathbf{P}_{\mathcal{U}} \mathbf{y}_2\|_F^2 \\
&= \min_{\mathcal{U} \in \mathbb{G}(n,r)} \left\| \begin{bmatrix} \mathbf{y}_1 - \mathbf{P}_{\mathcal{U}} \mathbf{P}_{\widehat{\mathcal{S}}_{b,r}} \mathbf{y}_1 & \mathbf{y}_2 - \mathbf{P}_{\mathcal{U}} \mathbf{y}_2 \end{bmatrix} \right\|_F^2 \\
&= \min_{\mathcal{U} \in \mathbb{G}(n,r)} \left\| \begin{bmatrix} \mathbf{P}_{\widehat{\mathcal{S}}_{b,r}^\perp} \mathbf{y}_1 & \mathbf{0}_{n \times b} \end{bmatrix} + \mathbf{P}_{\mathcal{U}^\perp} \begin{bmatrix} \mathbf{P}_{\widehat{\mathcal{S}}_{b,r}} \mathbf{y}_1 & \mathbf{y}_2 \end{bmatrix} \right\|_F^2 \\
&= \|\mathbf{P}_{\widehat{\mathcal{S}}_{b,r}^\perp} \mathbf{y}_1\|_F^2 + \min_{\mathcal{U} \in \mathbb{G}(n,r)} \left\| \mathbf{P}_{\mathcal{U}^\perp} \begin{bmatrix} \mathbf{P}_{\widehat{\mathcal{S}}_{b,r}} \mathbf{y}_1 & \mathbf{y}_2 \end{bmatrix} \right\|_F^2 \quad (\text{see the text below}) \\
&= \|\mathbf{P}_{\widehat{\mathcal{S}}_{b,r}^\perp} \mathbf{y}_1\|_F^2 + \min_{\mathcal{U} \in \mathbb{G}(n,r)} \left\| \mathbf{P}_{\mathcal{U}^\perp} \begin{bmatrix} \widehat{\mathbf{Y}}_{b,r} & \mathbf{y}_2 \end{bmatrix} \right\|_F^2, \quad (\text{see (18)}) \tag{21}
\end{aligned}$$

and let $\widehat{\mathcal{S}}_{2b,r}$ be a minimiser of the last program above. Above, \perp shows the orthogonal complement of a subspace. The second to last line above follows because $\widehat{\mathcal{S}}_{2b,r}$ is always within the column span of $[\mathbf{P}_{\widehat{\mathcal{S}}_{b,r}} \mathbf{y}_1 \ \mathbf{y}_2]$. Note also that $\widehat{\mathcal{S}}_{2b,r}$ is the span of leading r principal components of the matrix $[\widehat{\mathbf{Y}}_{1,r} \ \mathbf{y}_2]$, similar to Program (14). This again coincides with the output of MOSES in the second iteration, because

$$\begin{aligned}
\widehat{\mathbf{Y}}_{2b,r} &= \text{SVD}_r \left(\begin{bmatrix} \widehat{\mathbf{Y}}_{b,r} & \mathbf{y}_2 \end{bmatrix} \right) \quad (\text{see (8)}) \\
&= \mathbf{P}_{\widehat{\mathcal{S}}_{2b,r}} \begin{bmatrix} \widehat{\mathbf{Y}}_{b,r} & \mathbf{y}_2 \end{bmatrix}. \quad (\text{similar to the second line of (7)}) \tag{22}
\end{aligned}$$

Continuing this procedure precisely produces the iterates of MOSES. Therefore we might interpret MOSES as an optimisation algorithm for solving Program (14) by making a sequence of approximations.

4 Performance of MOSES

In this section, we study the performance of MOSES in a stochastic setup. Consider the probability space $(\mathbb{R}^n, \mathcal{B}, \mu)$, where \mathcal{B} is the Borel σ -algebra and μ is an *unknown* probability measure with zero mean, namely $\int_{\mathbb{R}^n} y \mu(dy) = 0$. We are interested in finding an r -dimensional subspace \mathcal{U} that best approximates the probability measure μ . That is, with y drawn from this probability space, we are interested in finding an r -dimensional subspace \mathcal{U} that minimises the *population risk*:

$$\begin{aligned}
\min_{\mathcal{U} \in \mathbb{G}(n,r)} \mathbb{E} \|y - \mathbf{P}_{\mathcal{U}} y\|_2^2 &= \min_{\mathcal{U} \in \mathbb{G}(n,r)} \int_{\mathbb{R}^n} \|y - \mathbf{P}_{\mathcal{U}} y\|_F^2 \mu(dy) \\
&=: \rho_r^2(\mu). \tag{23}
\end{aligned}$$

Since μ is unknown, we cannot directly solve Program (23) but suppose that instead we have access to the *training samples* $\{y_t\}_{t=1}^T \subset \mathbb{R}^n$ drawn independently from this probability space $(\mathbb{R}^n, \mathcal{B}, \mu)$. Let us form $\mathbf{Y}_T \in \mathbb{R}^{n \times T}$ by concatenating these vectors, see (1). In lieu of Program (23), we then replace the population risk above with the *empirical risk*:

$$\begin{aligned}
\min_{\mathcal{U} \in \mathbb{G}(n,r)} \frac{1}{T} \sum_{t=1}^T \|y_t - \mathbf{P}_{\mathcal{U}} y_t\|_2^2 &= \min_{\mathcal{U} \in \mathbb{G}(n,r)} \frac{1}{T} \|\mathbf{Y}_T - \mathbf{P}_{\mathcal{U}} \mathbf{Y}_T\|_F^2 \quad (\text{see (1)}) \\
&= \frac{1}{T} \|\mathbf{Y}_T - \mathbf{P}_{\mathcal{S}_{T,r}} \mathbf{Y}_T\|_F^2 \quad (\text{see the text below}) \\
&= \frac{1}{T} \|\mathbf{Y}_T - \mathbf{Y}_{T,r}\|_F^2 \quad (\mathbf{Y}_{T,r} = \text{SVD}_r(\mathbf{Y}_T)) \\
&=: \frac{\rho_r^2(\mathbf{Y}_T)}{T}. \quad (\text{see (5)}) \tag{24}
\end{aligned}$$

Here, $\mathcal{S}_{T,r} \in \mathbb{G}(n, r)$ is a minimiser of the above program with orthonormal basis $\mathbf{S}_{T,r} \in \mathbb{R}^{n \times r}$. Note that $\mathbf{S}_{T,r}$ consists of leading r principal components of \mathbf{Y}_T , namely it contains leading r left singular vectors of \mathbf{Y}_T by the Eckart-Young-Mirsky Theorem [19]. Given its principal components, we can then reduce the dimension of the data matrix $\mathbf{Y}_T \in \mathbb{R}^{n \times T}$ from n to r by computing $\mathbf{S}_{T,r}^* \mathbf{Y}_T \in \mathbb{R}^{r \times T}$. Note also that $\mathcal{S}_{T,r}$ is a possibly sub-optimal choice in Program (23), namely

$$\mathbb{E}_y \|y - \mathbf{P}_{\mathcal{S}_{T,r}} y\|_2^2 \geq \rho_r^2(\mu). \quad (\text{see (23)}) \quad (25)$$

But one would hope that \mathcal{S}_T still nearly minimises Program (23), in the sense that

$$\mathbb{E}_y \|y - \mathbf{P}_{\mathcal{S}_{T,r}} y\|_2^2 \approx \rho_r^2(\mu), \quad (26)$$

with high probability over the choice of training data $\{y_t\}_{t=1}^T$. That is, one would hope that the *generalisation error* of Program (24) is small. Above, \mathbb{E}_y stands for expectation over y , so that the left-hand side of (26) is still a random variable because of its dependence on \mathcal{S}_T .

If the training data $\{y_t\}_{t=1}^T$ is presented to us sequentially and little storage is available, we cannot hope to directly solve Program (24). Moreover, even if we have enough storage, we might not want to wait for all the data to arrive before solving Program (24). In this streaming scenario, we may apply MOSES to obtain the (rank- r) output $\widehat{\mathbf{Y}}_{T,r}$. We then set

$$\widehat{\mathcal{S}}_{T,r} = \text{span}(\widehat{\mathbf{Y}}_{T,r}), \quad (27)$$

with orthonormal basis $\widehat{\mathcal{S}}_T \in \mathbb{R}^{n \times r}$. Note that $\widehat{\mathcal{S}}_T$ is MOSES' estimate of leading r principal components of the data matrix \mathbf{Y}_T and is possibly suboptimal in the sense that

$$\|\mathbf{Y}_T - \widehat{\mathbf{Y}}_{T,r}\|_F \geq \rho_r(\mathbf{Y}_T). \quad (\text{see (24)}) \quad (28)$$

But we would still hope that the output $\widehat{\mathbf{Y}}_{T,r}$ of MOSES is a nearly optimal choice in Program (24), in the sense that

$$\|\mathbf{Y}_T - \widehat{\mathbf{Y}}_{T,r}\|_F \approx \rho_r(\mathbf{Y}_T), \quad (29)$$

with high probability over the choice of $\{y_t\}_{t=1}^T$. Moreover, as with (26), $\widehat{\mathcal{S}}_{T,r}$ is again a possibly suboptimal choice for Program (23), and yet we would hope that

$$\mathbb{E}_y \|y - \mathbf{P}_{\widehat{\mathcal{S}}_{T,r}} y\|_2^2 \approx \rho_r^2(\mu), \quad (30)$$

with high probability over the choice of $\{y_t\}_{t=1}^T$.

To summarise the discussion above, the key questions are whether (26,29,30) hold. Let us answer these questions for the important case where μ is a zero-mean Gaussian probability measure with covariance matrix $\mathbf{\Xi} \in \mathbb{R}^{n \times n}$. For this choice of μ in (23), it is not difficult to verify that

$$\rho_r^2(\mu) = \sum_{i=r+1}^n \lambda_i(\mathbf{\Xi}), \quad (31)$$

where $\lambda_1(\mathbf{\Xi}) \geq \lambda_2(\mathbf{\Xi}) \geq \dots$ are the eigenvalues of the covariance matrix $\mathbf{\Xi}$. From now on, let us use the shorthand

$$\rho_r = \rho_r(\mu), \quad \lambda_i = \lambda_i(\mathbf{\Xi}), \quad i \in [1 : n].$$

For our choice of μ above as a Gaussian measure with covariance matrix $\mathbf{\Xi} \in \mathbb{R}^{n \times n}$, one can use standard tools from the covariance estimation literature to show that (26) holds when T is sufficiently large, the proof of which is included in Appendix B of the supplementary material for completeness [21, 22, 23].

Proposition 1. *Suppose that $\{y_t\}_{t=1}^T \subset \mathbb{R}^n$ are drawn independently from a zero-mean Gaussian measure μ with covariance matrix $\mathbf{\Xi} \in \mathbb{R}^{n \times n}$ and form $\mathbf{Y}_T \in \mathbb{R}^{n \times T}$ by concatenating these vectors, see (1). Suppose also that $\mathcal{S}_{T,r} \in \mathbb{G}(n, r)$ is the span of leading r principal components of \mathbf{Y}_T . For $1 \leq \alpha \leq \sqrt{T/\log T}$, it then holds that*

$$\frac{\rho_r^2(\mathbf{Y}_T)}{T} \lesssim \alpha \rho_r^2, \quad (32)$$

$$\mathbb{E}_y \|y - \mathbf{P}_{\mathcal{S}_{T,r}} y\|_2^2 \lesssim \alpha \rho_r^2 + \alpha(n-r)\lambda_1 \sqrt{\frac{\log T}{T}}, \quad (33)$$

except with a probability of at most $T^{-C\alpha^2}$, where C is a universal constant. Here, \lesssim suppresses any universal constants for a more tidy presentation.

In words, (33) states that the generalisation error of Program (24) is small, namely (26) holds. Indeed, as T increases, the right-hand side of (33) approaches the residual squared of \mathbf{Y}_T/\sqrt{T} , which is bounded by $C\alpha\rho_r^2$. In particular, (26) holds when $\alpha = O(1)$ and T is sufficiently large. As the dimension r of the subspace that we fit to the data approaches the ambient dimension n , note that the right-hand side of (33) vanishes.

In contrast, MOSES operates in a streaming regime, where we are unable to fully store the data matrix \mathbf{Y}_T and consequently unable to find its principal components directly. That is, we cannot directly solve Program (24) in the streaming regime. However, Theorem 1 below states that MOSES approximately solves Program (24), namely MOSES approximately estimates the leading principal components of \mathbf{Y}_T and reduces the dimension of data from n to r with only $O(r(n+T))$ bits of memory, rather than $O(nT)$ bits required for solving Program (24) with ‘‘offline’’ truncated SVD. Moreover, MOSES approximately solves Program (23). In other words, MOSES satisfies *both* (29,30). These statements are made concrete below and proved in Appendix C of the supplementary material.

Theorem 1. (Performance of MOSES) *Suppose that $\{y_t\}_{t=1}^T \subset \mathbb{R}^n$ are drawn independently from a zero-mean Gaussian probability measure μ with covariance matrix $\Xi \in \mathbb{R}^{n \times n}$. Let us define*

$$\kappa_r^2 := \frac{\lambda_1}{\lambda_r}, \quad \rho_r^2 = \sum_{i=r+1}^n \lambda_i, \quad \eta_r := \kappa_r + \sqrt{\frac{2\alpha\rho_r^2}{p^{\frac{1}{3}}\lambda_r}}, \quad (34)$$

where $\lambda_1 \geq \lambda_2 \geq \dots$ are the eigenvalues of Ξ . Let $\widehat{\mathcal{S}}_{T,r} = \text{span}(\widehat{\mathbf{Y}}_{T,r})$ be the span of the output of MOSES, see (27). Then, for tuning parameters $1 \leq \alpha \leq \sqrt{T/\log T}$ and $p > 1$, it holds that

$$\frac{\|\mathbf{Y}_T - \widehat{\mathbf{Y}}_{T,r}\|_F^2}{T} \lesssim \frac{\alpha p^{\frac{1}{3}} 4^{p\eta_r^2}}{(p^{\frac{1}{3}} - 1)^2} \cdot \min(\kappa_r^2 \rho_r^2, r\lambda_1 + \rho_r^2) \left(\frac{T}{p\eta_r^2 b}\right)^{p\eta_r^2 - 1}, \quad (35)$$

$$\mathbb{E}_y \|y - \mathbf{P}_{\widehat{\mathcal{S}}_{T,r}} y\|_2^2 \lesssim \frac{\alpha p^{\frac{1}{3}} 4^{p\eta_r^2}}{(p^{\frac{1}{3}} - 1)^2} \cdot \min(\kappa_r^2 \rho_r^2, r\lambda_1 + \rho_r^2) \left(\frac{T}{p\eta_r^2 b}\right)^{p\eta_r^2 - 1} + \alpha(n-r)\lambda_1 \sqrt{\frac{\log T}{T}}, \quad (36)$$

except with a probability of at most $T^{-C\alpha^2} + e^{-C\alpha r}$ and provided that

$$b \geq \frac{\alpha p^{\frac{1}{3}} r}{(p^{\frac{1}{6}} - 1)^2}, \quad b \geq C\alpha r, \quad T \geq p\eta_r^2 b. \quad (37)$$

The requirement $T \geq p\eta_r^2 b$ in the last line above is only for a compact bound in (35,36) and is not necessary. A general expression for arbitrary T is given in the proof, see (72). A few remarks about Theorem 1 are in order.

Discussion of Theorem 1. On the one hand, Theorem 1 and specifically (35) state that (29) holds under certain conditions, namely MOSES approximately solves Program (14) or, in other words, MOSES successfully performs streaming (linear) dimensionality reduction. Indeed, (35) loosely speaking states that $\|\mathbf{Y}_T - \widehat{\mathbf{Y}}_{T,r}\|_F^2$ scales with $\rho_r^2 T^{p\eta_r^2} / b^{p\eta_r^2 - 1}$, whereas the residual squared of \mathbf{Y}_T scales with $\rho_r^2 T$ by (32). That

is,

$$\begin{aligned}\|\mathbf{Y}_T - \widehat{\mathbf{Y}}_{T,r}\|_F^2 &\propto \left(\frac{T}{b}\right)^{p\eta_r^2-1} \rho_r^2(\mathbf{Y}_T) \\ &= \left(\frac{T}{b}\right)^{p\eta_r^2-1} \|\mathbf{Y}_T - \mathbf{Y}_{T,r}\|_F^2, \quad (\text{see (24)})\end{aligned}\tag{38}$$

after ignoring the less important terms. In words, applying offline truncated SVD to \mathbf{Y}_T outperforms the streaming MOSES by a polynomial factor in T/b .

- This polynomial factor can be negligible when the covariance matrix Ξ of the Gaussian data distribution μ is well-conditioned ($\kappa_r = O(1)$) and has a small residual ($\rho_r^2 = O(\lambda_r)$), in which case we will have $\eta_r = O(1)$, see (34). With $p = O(1)$, (38) then reads as

$$\|\mathbf{Y}_T - \widehat{\mathbf{Y}}_{T,r}\|_F^2 \propto \left(\frac{T}{b}\right)^{O(1)} \rho_r^2(\mathbf{Y}_T).$$

In particular, when the covariance matrix of the data distribution is rank- r , we have by (31) that $\rho_r = 0$. Consequently, (38) reads as $\widehat{\mathbf{Y}}_{T,r} = \mathbf{Y}_{T,r} = \mathbf{Y}_T$, namely the outputs of offline truncated SVD and MOSES coincide. In fact, MOSES correctly identifies the r -dimensional span of incoming data after processing the very first block.

- At the cost of a larger multiplicative factor on the right-hand side of (35), one might reduce the power of T in the first term of (35) by choosing p closer to one.
- The dependence of our results on the condition number κ_r and residual ρ_r is very likely *not* an artifact of the proof techniques, see (34). Indeed, when $\kappa_r \gg 1$, certain directions are less often observed in the incoming data vectors $\{y_t\}_{t=1}^T$, which tilts the estimate of MOSES towards the dominant principal components corresponding to the very large singular values. Moreover, if $\rho_r \gg 1$, there are too many significant principal components and MOSES can at most “remember” r of them from its previous iteration. In this scenario, approximating the incoming data with a rank- r subspace is not a good idea in the first place, in the sense that the residual $\rho_r(\mathbf{Y}_T)$ corresponding to the offline truncated SVD will be large too, and we should perhaps increase the dimension r of the subspace that we wish to fit to the incoming data $\{y_t\}_{t=1}^T$.
- Note also that, as b increases, performance of MOSES naturally approaches that of the offline truncated SVD. In particular, when $b = T$, MOSES processes all of the data at once and reduces to offline truncated SVD. This trend is somewhat imperfectly reflected in (35).

On the other hand, Theorem 1 and specifically (36) state that (30) holds under certain conditions. Indeed, for sufficiently large T , (36) loosely speaking reads as

$$\begin{aligned}\mathbb{E}_y \|y - \mathbf{P}_{\widehat{\mathcal{S}}_{T,r}} y\|_2^2 &\propto \left(\frac{T}{b}\right)^{p\eta_r^2-1} \rho_r^2 \\ &= \left(\frac{T}{b}\right)^{p\eta_r^2-1} \min_{\mathcal{U} \in \mathbb{G}(n,r)} \mathbb{E} \|y - \mathbf{P}_{\mathcal{U}} y\|_2^2. \quad (\text{see Program (23)})\end{aligned}\tag{39}$$

That is, the output of MOSES is sub-optimal for Program (23) by a polynomial factor in T , which is negligible if the covariance matrix Ξ of the data distribution μ is well-conditioned and has a small residual, see the discussion above.

Spiked covariance model. A popular model in the statistics literature is the spiked covariance model, where the data vectors $\{y_t\}_{t=1}^T$ are drawn from a distribution with a covariance matrix Ξ . Under this model, Ξ is a low-rank perturbation of the identity matrix [18, 23], namely $\lambda_1(\Xi) = \dots = \lambda_r(\Xi) = \lambda$ and

$\lambda_{r+1}(\Xi) = \dots = \lambda_n(\Xi) = 1$. Proposition 1 in this case reads as

$$\mathbb{E}\|y - \mathbf{P}_{\mathcal{S}_{T,r}} y\|_2^2 \propto (n-r) + (n-r)\lambda\sqrt{\frac{\log T}{T}}, \quad (40)$$

where $\mathcal{S}_{T,r}$ spans leading r principal components of the data matrix \mathbf{Y}_T . In contrast, Theorem 1 roughly speaking states that

$$\mathbb{E}\|y - \mathbf{P}_{\widehat{\mathcal{S}}_{T,r}} y\|_2^2 \propto (n-r) \left(\frac{T\lambda}{bn}\right)^{\frac{n}{\lambda}} + (n-r)\lambda\sqrt{\frac{\log T}{T}}, \quad (41)$$

where $\widehat{\mathcal{S}}_{T,r}$ spans the output of MOSES. When $\lambda \gtrsim n \log(T/b) = n \log K$ in particular, we find that the error bounds in (40,41) are of the same order. That is, under the spiked covariance model, MOSES for streaming truncated SVD matches the performance of “offline” truncated SVD, provided that the underlying distribution has a sufficiently large spectral gap. In practice, (41) is often a conservative bound.

Proof strategy. Starting with (36), the proof of Theorem 1 in Appendix C of the supplementary material breaks down the error associated with MOSES into two components as

$$\mathbb{E}_y \|y - \mathbf{P}_{\widehat{\mathcal{S}}_{T,r}} y\|_2 \leq \frac{1}{T} \|\mathbf{Y}_T - \mathbf{P}_{\widehat{\mathcal{S}}_{T,r}} \mathbf{Y}_T\|_F^2 + \left| \frac{1}{T} \|\mathbf{Y}_T - \mathbf{P}_{\widehat{\mathcal{S}}_{T,r}} \mathbf{Y}_T\|_F^2 - \mathbb{E}_y \|y - \mathbf{P}_{\widehat{\mathcal{S}}_{T,r}} y\|_2^2 \right|. \quad (42)$$

That is, we bound the population risk with the empirical risk. We control the empirical risk in the first part of the proof by noting that

$$\begin{aligned} \|\mathbf{Y}_T - \mathbf{P}_{\widehat{\mathcal{S}}_{T,r}} \mathbf{Y}_T\|_F &= \|\mathbf{P}_{\widehat{\mathcal{S}}_{T,r}^\perp} \mathbf{Y}_T\|_F \\ &= \|\mathbf{P}_{\widehat{\mathcal{S}}_{T,r}^\perp} (\mathbf{Y}_T - \widehat{\mathbf{Y}}_{T,r})\|_F \quad (\text{see (27)}) \\ &\leq \|\mathbf{Y}_T - \widehat{\mathbf{Y}}_{T,r}\|_F, \end{aligned} \quad (43)$$

where the last line gauges how well the output of MOSES approximates the data matrix \mathbf{Y}_T , see (35). We then bound $\|\mathbf{Y}_T - \widehat{\mathbf{Y}}_{T,r}\|_F$ in two steps: As it is common in these types of arguments, the first step finds a deterministic upper bound for this norm, which is then evaluated for our particular stochastic setup.

- The deterministic bound appears in Lemma 2 and gives an upper bound for $\|\mathbf{Y}_T - \widehat{\mathbf{Y}}_{T,r}\|_F$ in terms of the overall “innovation”. Loosely speaking, the innovation $\|\mathbf{P}_{\mathcal{S}_{(k-1)b,r}^\perp} \mathbf{y}_k\|_F$ at iteration k is the part of the new data block \mathbf{y}_k that cannot be described by the leading r principal components of data arrived so far, which span the subspace $\mathcal{S}_{(k-1)b,r}$.
- The stochastic bound is given in Lemma 3 and uses a tight perturbation result.

Our argument so far yields an upper bound on the empirical loss $\|\mathbf{Y}_T - \mathbf{P}_{\widehat{\mathcal{S}}_{T,r}} \mathbf{Y}_T\|_F$ that holds with high probability. In light of (42), it remains to control

$$\begin{aligned} \left| \frac{1}{T} \|\mathbf{Y}_T - \mathbf{P}_{\widehat{\mathcal{S}}_{T,r}} \mathbf{Y}_T\|_F^2 - \mathbb{E}_y \|y - \mathbf{P}_{\widehat{\mathcal{S}}_{T,r}} y\|_2^2 \right| &= \frac{1}{T} \left| \|\mathbf{Y}_T - \mathbf{P}_{\widehat{\mathcal{S}}_{T,r}} \mathbf{Y}_T\|_F^2 - \mathbb{E} \|\mathbf{Y}_T - \mathbf{P}_{\widehat{\mathcal{S}}_{T,r}} \mathbf{Y}_T\|_F^2 \right| \\ &= \frac{1}{T} \left| \|\mathbf{P}_{\widehat{\mathcal{S}}_{T,r}^\perp} \mathbf{Y}_T\|_F^2 - \mathbb{E} \|\mathbf{P}_{\widehat{\mathcal{S}}_{T,r}^\perp} \mathbf{Y}_T\|_F^2 \right| \end{aligned} \quad (44)$$

with a standard large deviation bound.

Other stochastic models. While our results were restricted to the Gaussian distribution, they extend easily and with minimal change to the larger class of subgaussian distributions. Beyond subgaussian data models, Lemma 2 is the key deterministic result, relating the MOSES error to the overall innovation. One might therefore control the overall innovation, namely the right-hand side of (69) in Lemma 2, for any other stochastic model at hand.

5 Prior Art

In this paper, we presented MOSES for streaming (linear) dimensionality reduction, an algorithm with minimal storage and computational requirements. One might think of MOSES as an online “subspace tracking” algorithm that identifies the linear structure of data as it arrives. Once the data has fully arrived, both principal components and the projected data are readily made available by MOSES and the user could immediately proceed with any additional learning and inference tasks. Note also that t in our notation need not correspond to time, see (1). For example, only a small portion of a large data matrix \mathbf{Y}_T can be stored in the fast access memory of the processing unit, which could instead use MOSES to fetch and process the data in small chunks and iteratively update its estimate of leading principal components. Moreover, MOSES can be easily adapted to the *dynamic* case where the distribution of data changes over time. In dynamic subspace tracking, each data vector y_t is drawn from a subspace $\mathcal{S}(t) \in \mathbb{G}(n, r)$ that might vary with time.

A closely related line of work is the incremental SVD [10, 11, 12, 13, 14]. Incremental SVD is a streaming algorithm that, given the (truncated) SVD of $\mathbf{Y}_{t-1} \in \mathbb{R}^{n \times (t-1)}$, aims to compute the (truncated) SVD of $\mathbf{Y}_t = [\mathbf{Y}_{t-1} \ y_t] \in \mathbb{R}^{n \times t}$, where $y_t \in \mathbb{R}^n$ is the newly arrived data vector and \mathbf{Y}_{t-1} is the matrix formed by concatenating the previous data vectors, see (1). It is easy to verify that MOSES slightly generalises incremental SVD to handle data blocks, see Algorithm 1. This small difference between incremental SVD and MOSES is in part what enables us to complement MOSES with a comprehensive statistical analysis in Theorem 1 which is, to the best of our knowledge, not available for incremental SVD, despite its long history and popularity. Indeed, [15] only very recently provided stochastic analysis for two of the variants of incremental SVD in [16, 17]. The results in [15] hold in expectation and for the special case of $r = 1$, the first leading principal component. Crucially, these results measure the angle $\angle[\mathcal{S}_{T,r}, \widehat{\mathcal{S}}_{T,r}]$ between the true leading principal components of the data matrix and those estimated by incremental SVD. In this sense, these types of results are inconclusive because incremental SVD estimates both left and right leading singular vectors of the data matrix, namely incremental SVD both estimates the leading principal components of the data matrix $\widehat{\mathcal{S}}_{T,r}$ and reduces the dimension of data by computing $\widehat{\mathcal{S}}_{T,r}^* \widehat{\mathbf{Y}}_{T,r} \in \mathbb{R}^{r \times T}$, where $\widehat{\mathbf{Y}}_{T,r}$ is the final output of incremental SVD. In contrast to [15], Theorem 1 and specifically (35) assesses the quality of both of these tasks and establishes that, under certain conditions, MOSES performs nearly as well as offline SVD. GROUSE [24] is a closely related algorithm for streaming PCA (on data with possibly missing entries) that can be interpreted as projected stochastic gradient descent on the Grassmannian manifold. GROUSE is effectively identical to incremental SVD when the incoming data is low-rank [24]. In [25], the authors offer theoretical guarantees for GROUSE that again does not account for the projected data and are based on the proof techniques of [15]. Their results hold without any missing data, in expectation, and in a setup similar to the spiked covariance model. An alternative to GROUSE is SNIPE that has much stronger theoretical guarantees in case of missing data [26, 27]. In Section 6, we numerically compare MOSES with GROUSE.

One might also view MOSES as a stochastic algorithm for PCA. Indeed, note that Program (23) is equivalent to

$$\begin{cases} \max_{\mathbf{U}^* \mathbf{U} = \mathbf{I}_r} \mathbb{E}_y \|\mathbf{U} \mathbf{U}^* y\|_F^2 \\ \mathbf{U}^* \mathbf{U} = \mathbf{I}_r \end{cases} = \begin{cases} \max_{\mathbf{U}^* \mathbf{U} = \mathbf{I}_r} \mathbb{E}_y \langle \mathbf{U} \mathbf{U}^*, yy^* \rangle \\ \mathbf{U}^* \mathbf{U} \preceq \mathbf{I}_r \end{cases} = \begin{cases} \max_{\mathbf{U}^* \mathbf{U} \preceq \mathbf{I}_r} \mathbb{E}_y \langle \mathbf{U} \mathbf{U}^*, yy^* \rangle \\ \mathbf{U}^* \mathbf{U} \preceq \mathbf{I}_r \end{cases} \quad (45)$$

where the maximisation is over matrix $\mathbf{U} \in \mathbb{R}^{n \times r}$. Above, $\mathbf{U}^* \mathbf{U} \preceq \mathbf{I}_r$ is the unit ball with respect to the spectral norm and $\mathbf{A} \preceq \mathbf{B}$ means that $\mathbf{B} - \mathbf{A}$ is a positive semi-definite matrix. The last identity above holds because a convex function is always maximised on the boundary of the feasible set. Using the Schur complement, we can equivalently write the last program above as

$$\begin{cases} \max_{\begin{bmatrix} \mathbf{I}_n & \mathbf{U} \\ \mathbf{U}^* & \mathbf{I}_r \end{bmatrix} \succcurlyeq \mathbf{0}_{(n+r) \times (n+r)}} \mathbb{E} \langle \mathbf{U} \mathbf{U}^*, yy^* \rangle \\ \begin{bmatrix} \mathbf{I}_n & \mathbf{U} \\ \mathbf{U}^* & \mathbf{I}_r \end{bmatrix} \succcurlyeq \mathbf{0}_{(n+r) \times (n+r)} \end{cases} = \begin{cases} \max_{\begin{bmatrix} \mathbf{I}_n & \mathbf{U} \\ \mathbf{U}^* & \mathbf{I}_r \end{bmatrix} \succcurlyeq \mathbf{0}_{(n+r) \times (n+r)}} \langle \mathbf{U} \mathbf{U}^*, \mathbf{\Xi} \rangle \\ \begin{bmatrix} \mathbf{I}_n & \mathbf{U} \\ \mathbf{U}^* & \mathbf{I}_r \end{bmatrix} \succcurlyeq \mathbf{0}_{(n+r) \times (n+r)} \end{cases}, \quad (46)$$

where $\mathbf{\Xi} = \mathbb{E}[yy^*] \in \mathbb{R}^{n \times n}$ is the covariance matrix of the data distribution μ . Note that Program (46) has a convex (in fact, quadratic) objective function that is *maximised* on a convex (conic) feasible set. We cannot hope to directly compute the gradient of the objective function above, namely $2\mathbf{\Xi}\mathbf{U}$, because the distribution

of y and hence its covariance matrix Ξ are unknown. Given an iterate $\widehat{\mathbf{S}}_t$, one might instead draw a random vector y_{t+1} from the probability measure μ and move along the direction $2y_{t+1}y_{t+1}^*\widehat{\mathbf{S}}_t$, motivated by the observation that $\mathbb{E}[2y_{t+1}y_{t+1}^*\widehat{\mathbf{S}}_t] = 2\Xi\widehat{\mathbf{S}}_t$. This is then followed by back projection onto the feasible set of Program (45). That is,

$$\widehat{\mathbf{S}}_{t+1} = \mathcal{P}\left(\mathbf{S}_t + 2\alpha_{t+1}y_{t+1}y_{t+1}^*\widehat{\mathbf{S}}_t\right), \quad (47)$$

for an appropriate step size α_{t+1} . Above, $\mathcal{P}(\mathbf{A})$ projects onto the unit spectral norm ball by setting to one all singular values of \mathbf{A} that exceed one. The stochastic projected gradient ascent for PCA, described above, is itself closely related to the so-called *power method* and is at the heart of [28, 29, 30, 31, 32], all lacking a statistical analysis similar to Theorem 1. One notable exception is the power method in [28] which in a sense applies *mini-batch* stochastic projected gradient ascent to solve Program (46), with data blocks (namely, batches) of size $b = \Omega(n)$. There the authors offer statistical guarantees for the spiked covariance model, see Section 4. As before, these guarantees are for the quality of estimated principal components and silent about the quality of projected data, which is addressed in Theorem 1. Note also that, especially when the data dimension n is large, one disadvantage of this approach is its large block size; it takes a long time of $\Omega(n)$ for the algorithm to update its estimate of the principal components. In this setup, we may think of MOSES as a stochastic algorithm for PCA based on alternative minimisation rather than gradient ascent, see Section 3. Moreover, MOSES updates its estimate frequently, after receiving every $b = O(r)$ data vectors, and also maintains the projected data. In Section 6, we numerically compare MOSES with the power method in [28]. A few closely related works are [33, 34, 35, 34].

In the context of online learning and *regret minimisation*, [36, 32] offer two algorithms the former of which is not memory optimal and the latter does not have guarantees similar to Theorem 1. See also [37]. A Bayesian approach to PCA is studied in [38, 39]. The *expectation maximisation* algorithm there could be implemented in an online fashion but without theoretical guarantees.

Ideas from sketching and randomised linear algebra could be integrated into MOSES and other streaming dimensionality reduction algorithms [40, 41, 42, 43, 44, 45]. It is also perhaps worth pointing out that one might consider a streaming algorithm as a special case of distributed computing along the “cone” tree shown in Figure 2, see also [46, 47]. When the data vectors have missing entries, a closely related problem is low-rank matrix completion [48, 49, 50].




Figure 2: Streaming problems may be interpreted as a special case of distributed computing. Each data block y_k lives on a node of the chain graph and the nodes are combined, from left to right, following the structure of the “cone” tree.

6 Experiments

In this section, we investigate the numerical performance of MOSES and compare it against two alternative algorithms, namely GROUSE [24] and the power method in [51], on both synthetic and real-world datasets. In all of our experiments, we reveal one by one the data vectors $\{y_t\}_{t=1}^T \subset \mathbb{R}^n$ and, for every t , hope to compute a rank- r truncated SVD of the data matrix arrived so far, namely $[y_1, \dots, y_t]$. For the tests that use synthetic data, the vectors $\{y_t\}_{t=1}^T$ are drawn independently from a zero-mean Gaussian distribution with covariance matrix $\Xi = \mathbf{S}\mathbf{A}\mathbf{S}^*$, where $\mathbf{S} \in \mathbb{R}^{n \times n}$ is a generic orthonormal basis obtained by orthogonalising a standard random Gaussian matrix. The entries of the diagonal matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$, namely the eigenvalues of the covariance matrix Ξ , are selected according to the power law: $\lambda_i = i^{-\alpha}$ for a positive α . To be more succinct, where possible we use MATLAB’s notation for specifying the value ranges in this section.

To assess the performance of MOSES, let $\mathbf{Y}_t = [y_1, \dots, y_t] \in \mathbb{R}^{n \times t}$ be the data received by time t and let $\widehat{\mathbf{Y}}_{t,r}^m$ be the output of MOSES at time t .¹ Then the error incurred by MOSES is

$$\frac{1}{t} \|\mathbf{Y}_t - \widehat{\mathbf{Y}}_{t,r}^m\|_F^2, \quad (48)$$

see Theorem 1. Recall from (5) that the above error is always worse (larger) than the residual of \mathbf{Y}_t , namely

$$\|\mathbf{Y}_t - \widehat{\mathbf{Y}}_{t,r}^m\|_F^2 \geq \|\mathbf{Y}_t - \mathbf{Y}_{t,r}\|_F^2 = \rho_r^2(\mathbf{Y}_t), \quad (\text{see (5)}) \quad (49)$$

where $\mathbf{Y}_{t,r} = \text{SVD}_r(\mathbf{Y}_t)$ is a rank- r truncated SVD of \mathbf{Y}_t and $\rho_r^2(\mathbf{Y}_t)$ is the corresponding residual. Later in this section, we compare MOSES against GROUSE [24] and power method [28], both described in Section 5. These algorithms only estimate the principal components of the data, as opposed to MOSES which also projects the data onto these estimates. More specifically, let $\widehat{\mathcal{S}}_{t,r}^g \in \mathbb{G}(n, r)$ and $\widehat{\mathcal{S}}_{t,r}^p \in \mathbb{G}(n, r)$ be the span of the output of GROUSE and the power method at time t , respectively. These algorithms then incur the errors

$$\frac{1}{t} \|\mathbf{Y}_t - \mathbf{P}_{\widehat{\mathcal{S}}_{t,r}^g} \mathbf{Y}_t\|_F^2, \quad \frac{1}{t} \|\mathbf{Y}_t - \mathbf{P}_{\widehat{\mathcal{S}}_{t,r}^p} \mathbf{Y}_t\|_F^2, \quad (50)$$

respectively. Above, $\mathbf{P}_{\mathcal{A}} \in \mathbb{R}^{n \times n}$ is the orthogonal projection onto the subspace \mathcal{A} . We now set out to do various tests and report the results. We remark that the accompanying MATLAB code is publicly available.²

Ambient dimension. On a synthetic dataset with $\alpha = 1$ and $T = 2000$, we first test MOSES by varying the ambient dimension as $n \in \{200 : 200 : 1200\}$, and setting the rank and block size to $r = 15$, $b = 2r = 30$. The average error over ten trials is reported in Figure 3a. Note that the error is increasing in n , see the discussion under spiked covariance model in Section 4.

Block size. On a synthetic dataset with $\alpha = 1$ and $T = 2000$, we test MOSES by setting the ambient dimension and rank to $n = 1200$, $r = 15$, and varying the block size as $b \in \{r : r : 15r\}$. The average error over ten trials is reported in Figure 3b. Note that the MOSES is robust against the choice of the block size and that, at the extreme case of $b = T$, error vanishes and MOSES reduces to “offline” truncated SVD.

Rank. On a synthetic dataset with $\alpha = 1$ and $T = 2000$, we test MOSES by setting the ambient dimension and block size to $n = 1200$, $b = 2r$, and varying the rank as $r \in \{5 : 5 : 25\}$. The average error over ten trials is reported in Figure 3c. As expected, the error is decreasing in the dimension r of the subspace that we fit to the data and in fact, at the extreme case of $r = n$, there would be no error at all.

Comparisons on synthetic datasets. On synthetic datasets with $\alpha \in \{0.01, 0.1, 0.5, 1\}$ and $T = 2000$, we compare MOSES against GROUSE and power method.³ More specifically, we set the ambient dimension

¹Note that MOSES updates its estimate after receiving each block of data, namely after every b data vectors. For the sake of an easier comparison with other algorithms (with different block sizes), we properly “interpolate” the outputs of all algorithms over time.

²<https://github.com/andylamp/moses>

³The MATLAB code for GROUSE is publicly available at <http://sunbeam.ece.wisc.edu/grouse>.




Figure 3: Performance of MOSES on synthetic datasets, see Section 6 for the details.

to $n = 200$ and the rank to $r = 10$. For MOSES, the block size was set to $b = 2r$. For GROUSE and power method, we set the step size and block size to 2 and $2n = 400$, respectively, as these values seemed to produced the best results overall. Both GROUSE and power method were initialised randomly, as prescribed in [24, 51]. The average errors of all three algorithms over ten trials versus time is shown in Figure 4. Because of its large blocks size of $O(n)$ [51], Note that the power method updates its estimate of principal components much slower than MOSES, but the two algorithms converge to similar errors. The slow updates of power method will become a problem when working with dynamic data, where the distribution of arriving data changes over time. We will also see later that MOSES is much faster than the power method.

In order to better evaluate the practicality of our method we also evaluate all these three algorithms on actual, publicly available datasets; we use four different datasets that contain *mote* (sensor node) voltage, humidity, light, and temperature measurements over time [52]. These datasets were selected because they are publicly available and are representative of real-world applications due to their ambient dimension n being sufficiently large (> 45) to reflect practical deployments.

Comparison on the mote voltage dataset. The first dataset we evaluate has an ambient dimension of $n = 46$ and has $T = 7712$ columns; it is an inherently volatile dataset as it contains the rapid small voltage changes the motes exhibit during operation. With the rest of parameters as described in the synthetic comparison above, the errors over time for all algorithms is shown in Figure 5a in logarithmic scale. MOSES here outperforms both GROUSE and power method.




Figure 4: Comparisons on synthetic datasets, see Section 6 for the details.

Comparison on the mote humidity dataset. The second dataset evaluated has an ambient dimension of $n = 48$ and has $T = 7712$ columns. This dataset contains the humidity measurements of motes and is more periodic in nature with a larger range than the voltage dataset. With the rest of parameters as described in the synthetic comparison above, the errors over time for all algorithms is shown in Figure 5b in logarithmic scale. MOSES again outperforms the other two algorithms.

Comparison on the mote light dataset. The third dataset has an ambient dimension $n = 48$ and has $T = 7712$ columns. This dataset contains the light measurements of the motes and contains much more frequent value changes while having the highest range of all four datasets studied in this section. With the rest of parameters as described in the synthetic comparison above, the errors over time for all algorithms is shown in Figure 5c in logarithmic scale. As before, MOSES outperforms the other two algorithms.

Comparison on the mote temperature dataset. The last real dataset we consider in this instance has an ambient dimension of $n = 56$ and has $T = 7712$ columns. This dataset contains the temperature measurements of the sensor motes and has mostly periodic value changes and infrequent spikes. With the rest of parameters as described in the synthetic comparison above, the errors over time for all algorithms is shown in Figure 5d in logarithmic scale. It is evident that MOSES outperforms the other two algorithms.

Computational complexity on synthetic datasets. Let us now turn our attention to the computational performance of these three algorithms. On synthetic datasets with $\alpha = 1$ and $T = 2000$, we compare the run-time of MOSES to GROUSE and power method, where the block sizes of MOSES and power method,




Figure 5: Comparisons on real-world datasets, see Section 6 for the details.

and the step size of GROUSE were set as described in the synthetic tests earlier. This simulation was carried out with MATLAB 2017b on a 2012 Mac Pro configured with Dual 6-core Intel Xeon X5690 with 64GB of DDR3 ECC RAM. The average run-time of all three algorithms over five trials and for various choices of rank r is shown in Figure 6. We note that the computational cost of MOSES remains consistently small throughout these simulations, especially for large ambient dimensions and ranks where GROUSE and power method perform poorly, see Figure 6c.

Acknowledgements

AE is supported by the Alan Turing Institute under the EPSRC grant EP/N510129/1 and also by the Turing Seed Funding grant SF019. RAH is supported by EPSRC grant EP/N510129/1. AG is supported by the Alan Turing Institute under the EPSRC grant EP/N510129/1 and TU/C/000003. AE is grateful to Chinmay Hedge, Mike Wakin, Jared Tanner, and Mark Davenport for insightful suggestions and valuable feedback. Parts of this project were completed when AE was a Leibniz Fellow at Oberwolfach Research Institute for Mathematics and AE is extremely grateful for their hospitality.

References

- [1] P. van Overschee and B. L. de Moor. *Subspace identification for linear systems: Theory, implementation, applications*. Springer US, 2012.
- [2] B. A. Ardekani, J. Kershaw, K. Kashikura, and I. Kanno. Activation detection in functional MRI using subspace modeling and maximum likelihood estimation. *IEEE Transactions on Medical Imaging*, 18(2):101–114, 1999.

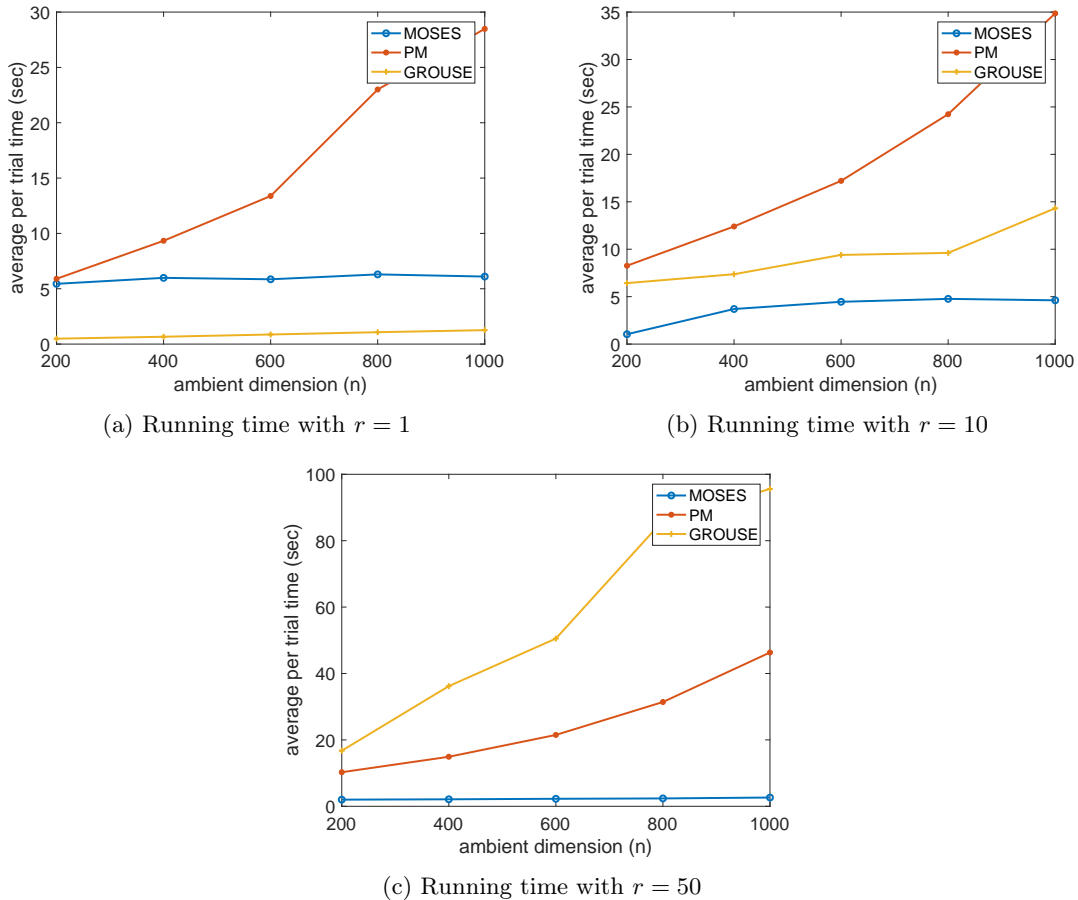


Figure 6: Computational complexity of all algorithms on synthetic datasets, see Section 6 for the details.

- [3] H. Krim and M. Viberg. Two decades of array signal processing research: The parametric approach. *IEEE Signal processing magazine*, 13(4):67–94, 1996.
- [4] L. Tong and S. Perreau. Multichannel blind identification: From subspace to maximum likelihood methods. *Proceedings of IEEE*, 86:1951–1968, 1998.
- [5] R. Vidal, Y. Ma, and S. Sastry. *Generalized Principal Component Analysis*. Interdisciplinary Applied Mathematics. Springer New York, 2016.
- [6] Raphael H Hauser and Armin Eftekhari. Pca by optimisation of symmetric functions has no spurious local optima. *arXiv preprint arXiv:1805.07459*, 2018.
- [7] Raphael A Hauser, Armin Eftekhari, and Heinrich F Matzinger. Pca by determinant optimization has no spurious local optima. *arXiv preprint arXiv:1803.04049*, 2018.
- [8] T. Hastie, R. Tibshirani, and J. Friedman. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer Series in Statistics. Springer New York, 2013.
- [9] L. Balzano, R. Nowak, and B. Recht. Online identification and tracking of subspaces from highly incomplete information. In *Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 704–711. IEEE, 2010.
- [10] James R Bunch, Christopher P Nielsen, and Danny C Sorensen. Rank-one modification of the symmetric eigenproblem. *Numerische Mathematik*, 31(1):31–48, 1978.
- [11] Matthew Brand. Fast low-rank modifications of the thin singular value decomposition. *Linear algebra and its applications*, 415(1):20–30, 2006.
- [12] Matthew Brand. Incremental singular value decomposition of uncertain data with missing values. *Computer Vision?ECCV 2002*, pages 707–720, 2002.
- [13] Pierre Comon and Gene H Golub. Tracking a few extreme singular values and vectors in signal processing. *Proceedings of the IEEE*, 78(8):1327–1343, 1990.
- [14] Yongmin Li. On incremental and robust subspace learning. *Pattern recognition*, 37(7):1509–1518, 2004.

- [15] A. Balsubramani, S. Dasgupta, and Y. Freund. The fast convergence of incremental pca. In *Advances in Neural Information Processing Systems*, pages 3174–3182, 2013.
- [16] TP Krasulina. The method of stochastic approximation for the determination of the least eigenvalue of a symmetrical matrix. *USSR Computational Mathematics and Mathematical Physics*, 9(6):189–195, 1969.
- [17] E. Oja. *Subspace methods of pattern recognition*. Electronic & electrical engineering research studies. Research Studies Press, 1983.
- [18] Iain M Johnstone. On the distribution of the largest eigenvalue in principal components analysis. *Annals of statistics*, pages 295–327, 2001.
- [19] C. Eckart and G. Young. The approximation of one matrix by another of lower rank. *Psychometrika*, 1:211–218, 1936.
- [20] L. Mirsky. Symmetric gauge functions and unitarily invariant norms. *Quart. J. Math. Oxford*, pages 1156–1159, 1966.
- [21] Armin Eftekhari, Ping Li, Michael B Wakin, and Rachel A Ward. Learning the differential correlation matrix of a smooth function from point samples. *arXiv preprint arXiv:1612.06339*, 2016.
- [22] Raphael Hauser, Raul Kangro, Jüri Lember, and Heinrich Matzinger. Quantifying the estimation error of principal components. *arXiv preprint arXiv:1710.10124*, 2017.
- [23] Roman Vershynin. How close is the sample covariance matrix to the actual covariance matrix? *Journal of Theoretical Probability*, 25(3):655–686, 2012.
- [24] L. Balzano and S. J Wright. On GROUSE and incremental SVD. In *IEEE International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP)*, pages 1–4. IEEE, 2013.
- [25] Dejiao Zhang and Laura Balzano. Global convergence of a grassmannian gradient descent algorithm for subspace estimation. In *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*, page 1460. [doi:10.14688](https://doi.org/10.14688), 2016.
- [26] Armin Eftekhari, Gregory Ongie, Laura Balzano, and Michael B Wakin. Streaming principal component analysis from incomplete data. *arXiv preprint arXiv:1612.00904*, 2018.
- [27] A. Eftekhari, L. Balzano, and M. B. Wakin. What to expect when you are expecting on the Grassmannian. *arXiv preprint arXiv:1611.07216*, 2016.
- [28] Ioannis Mitliagkas, Constantine Caramanis, and Prateek Jain. Memory limited, streaming pca. In *Advances in Neural Information Processing Systems*, pages 2886–2894, 2013.
- [29] E. Oja and J. Karhunen. On stochastic approximation of the eigenvectors and eigenvalues of the expectation of a random matrix. *Journal of mathematical analysis and applications*, 106(1):69–84, 1985.
- [30] Terence D Sanger. Optimal unsupervised learning in a single-layer linear feedforward neural network. *Neural networks*, 2(6):459–473, 1989.
- [31] Kwang In Kim, Matthias O Franz, and Bernhard Scholkopf. Iterative kernel principal component analysis for image modeling. *IEEE transactions on pattern analysis and machine intelligence*, 27(9):1351–1366, 2005.
- [32] Raman Arora, Andrew Cotter, Karen Livescu, and Nathan Srebro. Stochastic optimization for pca and pls. In *Communication, Control, and Computing (Allerton), 2012 50th Annual Allerton Conference on*, pages 861–868. IEEE, 2012.
- [33] Moritz Hardt and Eric Price. The noisy power method: A meta algorithm with applications. In *Advances in Neural Information Processing Systems*, pages 2861–2869, 2014.
- [34] Christopher De Sa, Kunle Olukotun, and Christopher Ré. Global convergence of stochastic gradient descent for some non-convex matrix problems. *arXiv preprint arXiv:1411.1134*, 2014.
- [35] Prateek Jain, Chi Jin, Sham M Kakade, Praneeth Netrapalli, and Aaron Sidford. Streaming pca: Matching matrix bernstein and near-optimal finite sample guarantees for oja’s algorithm. In *Conference on Learning Theory*, pages 1147–1164, 2016.
- [36] Manfred K Warmuth and Dima Kuzmin. Randomized online pca algorithms with regret bounds that are logarithmic in the dimension. *Journal of Machine Learning Research*, 9(Oct):2287–2320, 2008.
- [37] Christos Boutsidis, Dan Garber, Zohar Karnin, and Edo Liberty. Online principal components analysis. In *Proceedings of the twenty-sixth annual ACM-SIAM symposium on Discrete algorithms*, pages 887–901. Society for Industrial and Applied Mathematics, 2015.
- [38] Sam T Roweis. Em algorithms for pca and spca. In *Advances in neural information processing systems*, pages 626–632, 1998.
- [39] Michael E Tipping and Christopher M Bishop. Probabilistic principal component analysis. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 61(3):611–622, 1999.
- [40] Joel A Tropp, Alp Yurtsever, Madeleine Udell, and Volkan Cevher. Fixed-rank approximation of a positive-semidefinite matrix from streaming data. In *Advances in Neural Information Processing Systems*, pages 1225–1234, 2017.
- [41] Jiawei Chiu and Laurent Demanet. Sublinear randomized algorithms for skeleton decompositions. *SIAM Journal on Matrix Analysis and Applications*, 34(3):1361–1383, 2013.
- [42] Farhad Pourkamali-Anaraki and Stephen Becker. Randomized clustered nystrom for large-scale kernel machines. *arXiv preprint arXiv:1612.06470*, 2016.
- [43] Alex Gittens and Michael W Mahoney. Revisiting the nystrom method for improved large-scale machine learning. *The Journal of Machine Learning Research*, 17(1):3977–4041, 2016.

- [44] Mina Ghashami, Edo Liberty, Jeff M Phillips, and David P Woodruff. Frequent directions: Simple and deterministic matrix sketching. *SIAM Journal on Computing*, 45(5):1762–1792, 2016.
- [45] Anna C Gilbert, Jae Young Park, and Michael B Wakin. Sketched svd: Recovering spectral features from compressive measurements. *arXiv preprint arXiv:1211.0361*, 2012.
- [46] MA Iwen and BW Ong. A distributed and incremental svd algorithm for agglomerative data analysis on large networks. *SIAM Journal on Matrix Analysis and Applications*, 37(4):1699–1718, 2016.
- [47] Ahmed Sameh, Bernard Philippe, Dani Mezher, and Michael W Berry. Parallel algorithms for the singular value decomposition. In *Handbook of parallel computing and statistics*, pages 133–180. Chapman and Hall/CRC, 2005.
- [48] Mark A Davenport and Justin Romberg. An overview of low-rank matrix recovery from incomplete observations. *IEEE Journal of Selected Topics in Signal Processing*, 10(4):608–622, 2016.
- [49] Armin Eftekhari, Dehui Yang, and Michael B Wakin. Weighted matrix completion and recovery with prior subspace information. *IEEE Transactions on Information Theory*, 2018.
- [50] A. Eftekhari, M. B. Wakin, and R. A. Ward. MC²: A two-phase algorithm for leveraged matrix completion. *arXiv preprint arXiv:1609.01795*, 2016.
- [51] I. Mitliagkas, C. Caramanis, and P. Jain. Streaming PCA with many missing entries. *Preprint*, 2014.
- [52] Amol Deshpande, Carlos Guestrin, Samuel R Madden, Joseph M Hellerstein, and Wei Hong. Model-driven data acquisition in sensor networks. In *Proceedings of the Thirtieth international conference on Very large data bases-Volume 30*, pages 588–599. VLDB Endowment, 2004.
- [53] R. Vershynin. Introduction to the non-asymptotic analysis of random matrices. In Y. C. Eldar and G. Kutyniok, editors, *Compressed Sensing: Theory and Applications*, pages 95–110. Cambridge University Press, 2012.
- [54] Mark Rudelson, Roman Vershynin, et al. Hanson-wright inequality and sub-gaussian concentration. *Electronic Communications in Probability*, 18, 2013.
- [55] P. Wedin. Perturbation bounds in connection with singular value decomposition. *BIT Numerical Mathematics*, 12(1):99–111, 1972.
- [56] M. Ledoux and M. Talagrand. *Probability in Banach Spaces: Isoperimetry and Processes*. Classics in Mathematics. Springer Berlin Heidelberg, 2013.

A Notation and Toolbox

This section collects the notation and a number of useful results in one place for the convenience of the reader. We will always use bold letters for matrices and calligraphic letters for subspaces, for example matrix \mathbf{A} and subspace \mathcal{S} . In particular, $\mathbf{0}_{a \times b}$ denotes the $a \times b$ matrix of all zeros. For integers $a \leq b$, we use the convention that $[a : b] = \{a, \dots, b\}$. We will also use MATLAB’s matrix notation to represent rows, columns, and blocks of matrices, for example $\mathbf{A}[1 : r, :]$ is the restriction of matrix \mathbf{A} to its first r rows. Throughout, C is an absolute constant, the value of which might change in every appearance.

In the appendices, $\lambda_1(\mathbf{A}) \geq \lambda_2(\mathbf{A}) \geq \dots$ denote the eigenvalues of a symmetric matrix \mathbf{A} and $\sigma_1(\mathbf{B}) \geq \sigma_2(\mathbf{B}) \geq \dots$ denotes the singular values of a matrix \mathbf{B} . Also $\rho_r^2(\mathbf{B}) = \sum_{i \geq r+1} \sigma_i^2(\mathbf{B})$ stands for the residual of matrix \mathbf{B} .

Let us also recall some of the spectral properties of a standard random Gaussian matrix, namely a matrix populated with independent random Gaussian variables with zero-mean and unit variance. For a standard Gaussian matrix $\mathbf{G} \in \mathbb{R}^{a \times b}$ with $a \geq b$ and for fixed $\alpha \geq 1$, Corollary 5.35 in [53] dictates that

$$\sqrt{a} - \alpha\sqrt{b} \leq \sigma_b(\mathbf{G}) \leq \sigma_1(\mathbf{G}) \leq \sqrt{a} + \alpha\sqrt{b}, \quad (51)$$

except with a probability of at most $e^{-C\alpha^2 b}$. Moreover, for a matrix $\mathbf{\Gamma} \in \mathbb{R}^{a' \times a}$ and $\alpha \geq 1$, an application of the Hensen-Wright inequality [54, Theorem 1.1] yields that

$$\left| \|\mathbf{\Gamma}\mathbf{G}\|_F^2 - \mathbb{E}\|\mathbf{\Gamma}\mathbf{G}\|_F^2 \right| \leq \beta, \quad (52)$$

for $\beta \geq 0$ and except with a probability of at most

$$\exp\left(-\min\left(\frac{\beta^2}{b\|\mathbf{\Gamma}\|^2\|\mathbf{\Gamma}\|_F^2}, \frac{\beta}{\|\mathbf{\Gamma}\|^2}\right)\right),$$

where $\|\cdot\|$ stands for spectral norm. In particular, with the choice $\beta = \alpha^2 \|\mathbf{\Gamma}\|_F^2 b$ above and $\alpha \geq 1$, we find that

$$\|\mathbf{\Gamma}\mathbf{G}\|_F^2 \leq (1 + \alpha^2) \|\mathbf{\Gamma}\|_F^2 b \leq 2\alpha^2 \|\mathbf{\Gamma}\|_F^2 b, \quad (53)$$

except with a probability of at most

$$\exp(-C\alpha^2 b \|\mathbf{\Gamma}\|_F^2 / \|\mathbf{\Gamma}\|^2) \leq \exp(-C\alpha^2 b).$$

In a different regime, with the choice of $\beta = \alpha^2 \|\mathbf{\Gamma}\|_F^2 \sqrt{b}$ in (52) and $\alpha^2 \leq \sqrt{b}$, we arrive at

$$|\|\mathbf{\Gamma}\mathbf{G}\|_F^2 - \mathbb{E}\|\mathbf{\Gamma}\mathbf{G}\|_F^2| = |\|\mathbf{\Gamma}\mathbf{G}\|_F^2 - b\|\mathbf{\Gamma}\|_F^2| \leq \alpha^2 \|\mathbf{\Gamma}\|_F^2 \sqrt{b}, \quad (54)$$

except with a probability of at most

$$\exp(-C\alpha^4 \|\mathbf{\Gamma}\|_F^2 / \|\mathbf{\Gamma}\|^2) \leq \exp(-C\alpha^4).$$

B Proof of Proposition 1

Let

$$\mathbf{\Xi} = \mathbf{S}\mathbf{\Lambda}\mathbf{S}^* = \mathbf{S}\mathbf{\Sigma}^2\mathbf{S}^* \in \mathbb{R}^{n \times n} \quad (55)$$

be the eigen-decomposition of the covariance matrix $\mathbf{\Xi}$, where $\mathbf{S} \in \mathbb{R}^{n \times n}$ is an orthonormal matrix and the diagonal matrix $\mathbf{\Lambda} = \mathbf{\Sigma}^2 \in \mathbb{R}^{n \times n}$ contains the eigenvalues of $\mathbf{\Xi}$ in nonincreasing order, namely

$$\mathbf{\Lambda} = \mathbf{\Sigma}^2 = \begin{bmatrix} \sigma_1^2 & & & \\ & \sigma_2^2 & & \\ & & \ddots & \\ & & & \sigma_n^2 \end{bmatrix} \in \mathbb{R}^{n \times n}, \quad \sigma_1^2 \geq \sigma_2^2 \geq \dots \geq \sigma_n^2. \quad (56)$$

Throughout, we also make use of the condition number and residual, namely

$$\kappa_r = \frac{\sigma_1}{\sigma_r}, \quad \rho_r^2 = \rho_r^2(\mathbf{\Xi}) = \sum_{i=r+1}^n \sigma_i^2. \quad (\text{see (31)}) \quad (57)$$

Recall that $\{y_t\}_{t=1}^T \subset \mathbb{R}^n$ are the data vectors drawn from the Gaussian measure μ with zero mean and covariance matrix $\mathbf{\Xi}$, and that $\mathbf{Y}_T \in \mathbb{R}^{n \times T}$ is obtained by concatenating $\{y_t\}_{t=1}^T$. It follows that

$$y_t = \mathbf{S}\mathbf{\Sigma}g_t, \quad t \in [1 : T],$$

$$\mathbf{Y}_T = \mathbf{S}\mathbf{\Sigma}\mathbf{G}_T, \quad (58)$$

where $g_t \in \mathbb{R}^n$ and $\mathbf{G}_T \in \mathbb{R}^{n \times T}$ are standard random Gaussian vector and matrix, respectively. That is, g_t and \mathbf{G}_T are populated with independent Gaussian random variables with zero mean and unit variance. With these preparations, we are now ready to prove Proposition 1. For y drawn from the Gaussian measure

μ , note that

$$\begin{aligned}
\mathbb{E}_y \|y - \mathbf{P}_{S_{T,r}} y\|_2^2 &= \mathbb{E}_y \|\mathbf{P}_{S_{T,r}^\perp} y\|_2^2 \\
&= \mathbb{E}_y \langle \mathbf{P}_{S_{T,r}^\perp}, yy^* \rangle \\
&= \langle \mathbf{P}_{S_{T,r}^\perp}, \mathbf{\Xi} \rangle \\
&= \left\langle \mathbf{P}_{S_{T,r}^\perp}, \mathbf{\Xi} - \frac{\mathbf{Y}_T \mathbf{Y}_T^*}{T} \right\rangle + \frac{1}{T} \langle \mathbf{P}_{S_{T,r}^\perp}, \mathbf{Y}_T \mathbf{Y}_T^* \rangle \\
&= \left\langle \mathbf{P}_{S_{T,r}^\perp}, \mathbf{\Xi} - \frac{\mathbf{Y}_T \mathbf{Y}_T^*}{T} \right\rangle + \frac{1}{T} \|\mathbf{P}_{S_{T,r}^\perp} \mathbf{Y}_T\|_F^2 \\
&= \left\langle \mathbf{P}_{S_{T,r}^\perp}, \mathbf{\Xi} - \frac{\mathbf{Y}_T \mathbf{Y}_T^*}{T} \right\rangle + \frac{\rho_r^2(\mathbf{Y}_T)}{T} \quad (\text{see Program (24)}) \\
&= \frac{1}{T} \left(\mathbb{E} \|\mathbf{P}_{S_{T,r}^\perp} \mathbf{Y}_T\|_F^2 - \|\mathbf{P}_{S_{T,r}^\perp} \mathbf{Y}_T\|_F^2 \right) + \frac{\rho_r^2(\mathbf{Y}_T)}{T}. \quad (\text{see (58)}) \tag{59}
\end{aligned}$$

Let us next control the two components in the last line above. The first component above involves the deviation of random variable $\|\mathbf{P}_{S_{T,r}^\perp} \mathbf{Y}_T\|_F^2$ from its expectation. By invoking the Hensen-Wright inequality in Appendix A and for $\tilde{\alpha}^2 \leq \sqrt{T}$, we write that

$$\begin{aligned}
\mathbb{E} \|\mathbf{P}_{S_{T,r}^\perp} \mathbf{Y}_T\|_F^2 - \|\mathbf{P}_{S_{T,r}^\perp} \mathbf{Y}_T\|_F^2 &= \mathbb{E} \|\mathbf{P}_{S_{T,r}^\perp} \mathbf{S} \mathbf{\Sigma} \cdot \mathbf{G}_T\|_F^2 - \|\mathbf{P}_{S_{T,r}^\perp} \mathbf{S} \mathbf{\Sigma} \cdot \mathbf{G}_T\|_F^2 \quad (\text{see (58)}) \\
&\leq \tilde{\alpha}^2 \|\mathbf{P}_{S_{T,r}^\perp} \mathbf{S} \mathbf{\Sigma}\|_F^2 \sqrt{T} \quad (\text{see (54)}) \\
&\leq \tilde{\alpha}^2 \|\mathbf{P}_{S_{T,r}^\perp} \mathbf{S}\|_F^2 \|\mathbf{\Sigma}\|^2 \sqrt{T} \\
&\leq \tilde{\alpha}^2 (n-r) \sigma_1^2 \sqrt{T}, \quad (\text{see (56,57)}) \tag{60}
\end{aligned}$$

except with a probability of at most $e^{-C\tilde{\alpha}^4}$. In particular, for the choice of $\tilde{\alpha}^2 = \alpha^2 \sqrt{\log T}$ with $\alpha^2 \leq \sqrt{T/\log T}$, we find that

$$\mathbb{E} \|\mathbf{P}_{S_{T,r}^\perp} \mathbf{Y}_T\|_F^2 - \|\mathbf{P}_{S_{T,r}^\perp} \mathbf{Y}_T\|_F^2 \leq \alpha^2 (n-r) \sigma_1^2 \sqrt{T \log T}, \tag{61}$$

except with a probability of $T^{-C\alpha^4}$. We next bound the second term in the last line of (59), namely the residual of \mathbf{Y}_T . Note that

$$\begin{aligned}
\rho_r^2(\mathbf{Y}_T) &= \rho_r^2(\mathbf{S} \mathbf{\Sigma} \mathbf{G}_T) \quad (\text{see (58)}) \\
&= \rho_r^2(\mathbf{\Sigma} \mathbf{G}_T) \quad (\mathbf{S}^* \mathbf{S} = \mathbf{I}_n) \\
&= \min_{\text{rank}(\mathbf{X})=r} \|\mathbf{\Sigma} \mathbf{G}_T - \mathbf{X}\|_F^2. \quad (\text{see (57)}) \tag{62}
\end{aligned}$$

By substituting above the suboptimal choice of

$$\mathbf{X}_o = \begin{bmatrix} \mathbf{\Sigma}[1:r, 1:r] \cdot \mathbf{G}_T[1:r, :] \\ \mathbf{0}_{(n-r) \times T} \end{bmatrix}, \tag{63}$$

we find that

$$\begin{aligned}
\rho_r^2(\mathbf{Y}_T) &= \min_{\text{rank}(\mathbf{X})=r} \|\mathbf{\Sigma} \mathbf{G}_T - \mathbf{X}\|_F^2 \quad (\text{see (62)}) \\
&\leq \|\mathbf{\Sigma} \mathbf{G}_T - \mathbf{X}_o\|_F^2 \\
&= \|\mathbf{\Sigma}[r+1:n, r+1:n] \cdot \mathbf{G}_T[r+1:n, :]\|_F. \quad (\text{see (63)}) \tag{64}
\end{aligned}$$

Note that $\mathbf{G}_T[r+1:n, :] \in \mathbb{R}^{(n-r) \times T}$ is a standard Gaussian matrix. For $\alpha \geq 1$, an application of the Hensen-Wright inequality in Appendix A therefore implies that

$$\begin{aligned} \rho_r^2(\mathbf{Y}_T) &\leq \|\boldsymbol{\Sigma}[r+1:n, r+1:n] \cdot \mathbf{G}_T[r+1:n, :]\|_F^2 \quad (\text{see (64)}) \\ &\leq 2\alpha^2 \|\boldsymbol{\Sigma}[r+1:n, r+1:n]\|_F^2 T \quad (\text{see (53)}) \\ &= 2\alpha^2 \rho_r^2 T, \quad (\text{see (57)}) \end{aligned} \tag{65}$$

except with a probability of at most $e^{-C\alpha^2 T}$. We now substitute the bounds in (61) and (65) back into (59) to arrive at

$$\mathbb{E}\|y - \mathbf{P}_{\mathcal{S}_{T,r}} y\|_2^2 \leq \alpha^2 (n-r) \sigma_1^2 \sqrt{T \log T} + 2\alpha^2 \rho_r^2, \tag{66}$$

when $\alpha^2 \leq \sqrt{T/\log T}$ and except with a probability of at most

$$T^{-C\alpha^4} + e^{-C\alpha^2 T} \leq T^{-C\alpha^4}, \quad \left(\alpha^2 \leq \sqrt{T/\log T}\right)$$

where we have used the abuse of notation in which C is a universal constant that is allowed to change in every appearance. This completes the proof of Proposition 1.

C Proof of Theorem 1

In the rest of this paper, we slightly unburden the notation by using $\mathbf{Y}_k \in \mathbb{R}^{n \times kb}$ to denote \mathbf{Y}_{kb} . For example, we will use $\mathbf{Y}_K \in \mathbb{R}^{n \times T}$ instead of \mathbf{Y}_T because $T = Kb$. We also write $\widehat{\mathcal{S}}_{k,r}$ instead of $\widehat{\mathcal{S}}_{kb,r}$. As with the proof of Proposition 1, we argue that

$$\begin{aligned} \mathbb{E}_y \|y - \mathbf{P}_{\widehat{\mathcal{S}}_{K,r}} y\|_2^2 &\leq \frac{1}{T} \left(\mathbb{E} \|\mathbf{P}_{\widehat{\mathcal{S}}_{K,r}} \mathbf{Y}_T\|_F^2 - \|\mathbf{P}_{\widehat{\mathcal{S}}_{K,r}} \mathbf{Y}_T\|_F^2 \right) + \frac{1}{T} \|\mathbf{P}_{\widehat{\mathcal{S}}_{K,r}} \mathbf{Y}_K\|_F^2 \quad (\text{similar to (59)}) \\ &\leq \alpha^2 (n-r) \sigma_1^2 \sqrt{\frac{\log T}{T}} + \frac{1}{T} \|\mathbf{P}_{\widehat{\mathcal{S}}_{K,r}} \mathbf{Y}_K\|_F^2 \quad (\text{see (61)}) \\ &= \alpha^2 (n-r) \sigma_1^2 \sqrt{\frac{\log T}{T}} + \frac{1}{T} \|\mathbf{P}_{\widehat{\mathcal{S}}_{K,r}} (\mathbf{Y}_K - \widehat{\mathbf{Y}}_{K,r})\|_F^2 \quad (\text{see (27)}) \\ &\leq \alpha^2 (n-r) \sigma_1^2 \sqrt{\frac{\log T}{T}} + \frac{1}{T} \|\mathbf{Y}_K - \widehat{\mathbf{Y}}_{K,r}\|_F^2, \end{aligned} \tag{67}$$

except with a probability of at most $T^{-C\alpha^4}$ and provided that $\alpha^2 \leq \sqrt{T/\log T}$. It therefore remains to control the norm in the last line above. Recall that the output of MOSES, namely $\widehat{\mathbf{Y}}_{K,r}$, is intended to approximate a rank- r truncation of \mathbf{Y}_K . We will therefore compare the error $\|\mathbf{Y}_K - \widehat{\mathbf{Y}}_{K,r}\|_F$ in (67) with the true residual $\rho_r(\mathbf{Y}_K)$. To that end, our analysis consists of a deterministic bound and a stochastic evaluation of this bound. The deterministic bound is as follows, see Appendix D for the proof.

Lemma 2. *For every $k \in [1 : K]$, let $\mathbf{Y}_{k,r} = \text{SVD}_r(\mathbf{Y}_k) \in \mathbb{R}^{n \times kb}$ be a rank- r truncation of \mathbf{Y}_k and set $\mathcal{S}_{k,r} = \text{span}(\mathbf{Y}_{k,r}) \in \mathbb{G}(n, r)$. For $p > 1$, we also set*

$$\theta_k := 1 + \frac{p^{\frac{1}{3}} \|\mathbf{y}_k\|^2}{\sigma_r(\mathbf{Y}_{k-1})^2}. \tag{68}$$

Then the output of MOSES, namely $\widehat{\mathbf{Y}}_{K,r}$, satisfies

$$\|\mathbf{Y}_K - \widehat{\mathbf{Y}}_{K,r}\|_F^2 \leq \frac{p^{\frac{1}{3}}}{p^{\frac{1}{3}} - 1} \sum_{k=2}^K \left(\prod_{l=k+1}^K \theta_l \right) \|\mathbf{P}_{\mathcal{S}_{k-1,r}} \mathbf{y}_k\|_F^2, \tag{69}$$

where $\mathbf{P}_{\mathcal{S}_{k-1,r}} \in \mathbb{R}^{n \times n}$ is the orthogonal projection onto the orthogonal complement of $\mathcal{S}_{k-1,r}$. Above, we use

the convention that $\prod_{l=K+1}^K \theta_l = 1$.

In words, (69) gives a deterministic bound on the performance of MOSES. The term $\|\mathbf{P}_{\mathcal{S}_{k-1,r}^\perp} \mathbf{y}_k\|_F$ in (69) is in a sense the ‘‘innovation’’ at iteration k , namely the part of the new data block \mathbf{y}_k that cannot be described by the current estimate $\mathcal{S}_{k-1,r}$. The overall innovation in (69) clearly controls the performance of MOSES. In particular, if the data blocks are drawn from the same distribution, this innovation gradually reduces as k increases. For example, if $\{\mathbf{y}_k\}_{k=1}^K$ are drawn from a distribution with a rank- r covariance matrix, then the innovation term vanishes almost surely after finitely many iterations. In contrast, when the underlying covariance matrix is high-rank, the innovation term decays more slowly and never completely disappears even as $k \rightarrow \infty$. We will next evaluate the right-hand side of (69) in a stochastic setup, see Appendix G for the proof.

Lemma 3. *Suppose that $\{y_t\}_{t=1}^T$ are drawn from a zero-mean Gaussian probability measure with the covariance matrix $\Xi \in \mathbb{R}^{n \times n}$. Let $\sigma_1^2 \geq \sigma_2^2 \geq \dots$ be the eigenvalues of Ξ and recall the notation in (57). For $p > 1$, also let*

$$\eta_r := \kappa_r + \frac{\sqrt{2}\alpha\rho_r}{p^{\frac{1}{6}}\sigma_r}.$$

For $\alpha \geq 1$, it then holds that

$$\|\mathbf{Y}_K - \widehat{\mathbf{Y}}_{K,r}\|_F^2 \leq \frac{50p^{\frac{4}{3}}\alpha^2}{(p^{\frac{1}{3}} - 1)^2} \cdot \min(\kappa_r^2\rho_r^2, r\sigma_1^2 + \rho_r^2) \eta_r^2 b \left(\frac{2K}{p\eta_r^2} + 2\right)^{p\eta_r^2}, \quad (70)$$

except with a probability of at most $e^{-C\alpha^2 r}$ and provided that

$$b \geq \frac{p^{\frac{1}{3}}\alpha^2 r}{(p^{\frac{1}{6}} - 1)^2}, \quad b \geq C\alpha^2 r.$$

Substituting the right-hand side of (70) back into (67) yields that

$$\begin{aligned} \mathbb{E}_y \|y - \mathbf{P}_{\widehat{\mathcal{S}}_{K,r}} y\|_2^2 &\leq \alpha^2(n-r)\sigma_1^2 \sqrt{\frac{\log T}{T}} + \frac{1}{T} \|\mathbf{Y}_K - \widehat{\mathbf{Y}}_{K,r}\|_F^2, \quad (\text{see (67)}) \\ &\leq \alpha^2(n-r)\sigma_1^2 \sqrt{\frac{\log T}{T}} + \frac{50p^{\frac{4}{3}}\alpha^2}{(p^{\frac{1}{3}} - 1)^2} \cdot \min(\kappa_r^2\rho_r^2, r\sigma_1^2 + \rho_r^2) \frac{\eta_r^2}{K} \left(\frac{2K}{p\eta_r^2} + 2\right)^{p\eta_r^2}. \end{aligned} \quad (71)$$

In particular, if $K \geq p\eta_r^2$, we may simplify the above bound to read

$$\mathbb{E}_y \|y - \mathbf{P}_{\widehat{\mathcal{S}}_{K,r}} y\|_2^2 \leq \alpha^2(n-r)\sigma_1^2 \sqrt{\frac{\log T}{T}} + \frac{50p^{\frac{1}{3}}\alpha^2 4^{p\eta_r^2}}{(p^{\frac{1}{3}} - 1)^2} \cdot \min(\kappa_r^2\rho_r^2, r\sigma_1^2 + \rho_r^2) \left(\frac{K}{p\eta_r^2}\right)^{p\eta_r^2 - 1}, \quad (72)$$

which completes the proof of Theorem 1.

D Proof of Lemma 2

Recall that the output of MOSES is the sequence of rank- r matrices $\{\widehat{\mathbf{Y}}_k\}_{k=1}^K$. For every $k < K$, it is more convenient in the proof of Lemma 2 to pad both $\mathbf{Y}_k, \widehat{\mathbf{Y}}_{k,r} \in \mathbb{R}^{n \times kb}$ with zeros to form the $n \times Kb$ matrices

$$\left[\mathbf{Y}_k \quad \mathbf{0}_{n \times (K-k)b} \right], \quad \left[\widehat{\mathbf{Y}}_{k,r} \quad \mathbf{0}_{n \times (K-k)b} \right]. \quad (73)$$

We overload the notation $\mathbf{Y}_k, \widehat{\mathbf{Y}}_{k,r}$ to show the new $n \times Kb$ matrices in (73). Let

$$\begin{aligned} \widehat{\mathcal{S}}_{k,r} &= \text{span}(\widehat{\mathbf{Y}}_{k,r}) \in \mathbb{G}(n, r), \\ \widehat{\mathcal{Q}}_{k,r} &= \text{span}(\widehat{\mathbf{Y}}_{k,r}^*) \in \mathbb{G}(Kb, r) \end{aligned} \quad (74)$$

denote the (r -dimensional) column and row spaces of the rank- r matrix $\widehat{\mathbf{Y}}_{k,r} \in \mathbb{R}^{n \times Kb}$, respectively. Let also $\widehat{\mathbf{S}}_{k,r} \in \mathbb{R}^{n \times r}$ and $\widehat{\mathbf{Q}}_{k,r} \in \mathbb{R}^{Kb \times r}$ be orthonormal bases for these subspaces. We also let $\mathcal{I}_k \subset \mathbb{R}^{Kb}$ denote the b -dimensional subspace spanned by the coordinates $[(k-1)b+1 : kb]$, namely

$$\mathcal{I}_k = \text{span} \left(\begin{bmatrix} \mathbf{0}_{(k-1)b \times b} \\ \mathbf{I}_b \\ \mathbf{0}_{(K-k)b \times b} \end{bmatrix} \right) \in \mathbb{G}(Kb, b), \quad (75)$$

and we use the notation

$$\mathcal{J}_k := \mathcal{I}_1 \oplus \mathcal{I}_2 \cdots \oplus \mathcal{I}_k \in \mathbb{G}(Kb, kb), \quad k \in [1 : K], \quad (76)$$

to denote the kb -dimensional subspace that spans the first kb coordinates in \mathbb{R}^{Kb} . The following technical lemma, proved in Appendix E, gives another way of expressing the output of MOSES, namely $\{\widehat{\mathbf{Y}}_{k,r}\}_{k=1}^K$.

Lemma 4. *For every $k \in [1 : K]$, it holds that*

$$\widehat{\mathbf{Y}}_{k,r} = \mathbf{Y}_K \mathbf{P}_{\widehat{\mathbf{Q}}_{k,r}}, \quad (77)$$

or equivalently

$$\widehat{\mathbf{Y}}_{k-1,r} + \mathbf{Y}_k \mathbf{P}_{\mathcal{I}_k} = \mathbf{Y}_K \mathbf{P}_{\widetilde{\mathbf{Q}}_k}, \quad (78)$$

where

$$\widetilde{\mathbf{Q}}_k := \widehat{\mathbf{Q}}_{k-1,r} \oplus \mathcal{I}_k \subset \mathbb{R}^{Kb} \quad (79)$$

is the direct sum of the two subspaces $\widehat{\mathbf{Q}}_{k-1,r}$ and \mathcal{I}_k . In particular, the update rule (8) can be written as

$$\mathbf{Y}_K \mathbf{P}_{\widehat{\mathbf{Q}}_{k,r}} = \text{SVD}_r \left(\mathbf{Y}_K \mathbf{P}_{\widetilde{\mathbf{Q}}_k} \right), \quad k \in [2 : K]. \quad (80)$$

Lastly we have the inclusion

$$\widehat{\mathbf{Q}}_{k,r} \subset \widetilde{\mathbf{Q}}_k \subset \mathcal{J}_k \in \mathbb{G}(Kb, kb). \quad (81)$$

In particular, (77) and (81) together imply that

$$\widehat{\mathbf{Y}}_{k,r} = \mathbf{Y}_K \mathbf{P}_{\mathcal{J}_k} \mathbf{P}_{\widehat{\mathbf{Q}}_{k,r}} = \mathbf{Y}_k \mathbf{P}_{\widehat{\mathbf{Q}}_{k,r}},$$

that is, only \mathbf{Y}_k (containing the first kb data vectors) contributes to the formation of $\widehat{\mathbf{Y}}_{k,r}$, the output of algorithm at iteration k , which was to be expected of course. Recall that $\widehat{\mathbf{Y}}_{k,r}$ is intended to approximate $\mathbf{Y}_{k,r} = \text{SVD}_r(\mathbf{Y}_k)$. In light of Lemma 4, let us now derive a simple recursive expression for the residual $\mathbf{Y}_k - \widehat{\mathbf{Y}}_{k,r}$. For every $k \in [2 : K]$, it holds that

$$\begin{aligned} \mathbf{Y}_k - \widehat{\mathbf{Y}}_{k,r} &= \mathbf{Y}_K \mathbf{P}_{\mathcal{J}_k} - \mathbf{Y}_K \mathbf{P}_{\widehat{\mathbf{Q}}_{k,r}} \quad (\text{see (76) and (77)}) \\ &= \mathbf{Y}_K \mathbf{P}_{\mathcal{J}_{k-1}} + \mathbf{Y}_k \mathbf{P}_{\mathcal{I}_k} - \mathbf{Y}_K \mathbf{P}_{\widehat{\mathbf{Q}}_{k,r}} \quad (\text{see (76)}) \\ &= \mathbf{Y}_{k-1} + \mathbf{Y}_k \mathbf{P}_{\mathcal{I}_k} - \mathbf{Y}_K \mathbf{P}_{\widehat{\mathbf{Q}}_{k,r}} \quad (\text{see (76)}) \\ &= \mathbf{Y}_{k-1} - \widehat{\mathbf{Y}}_{k-1,r} + \mathbf{Y}_K \mathbf{P}_{\widehat{\mathbf{Q}}_{k-1,r}} + \mathbf{Y}_k \mathbf{P}_{\mathcal{I}_k} - \mathbf{Y}_K \mathbf{P}_{\widehat{\mathbf{Q}}_{k,r}} \quad (\text{see (77)}) \\ &= \left(\mathbf{Y}_{k-1} - \widehat{\mathbf{Y}}_{k-1,r} \right) + \mathbf{Y}_K \left(\mathbf{P}_{\widehat{\mathbf{Q}}_{k-1,r}} + \mathbf{P}_{\mathcal{I}_k} \right) - \mathbf{Y}_K \mathbf{P}_{\widehat{\mathbf{Q}}_{k,r}} \\ &= \left(\mathbf{Y}_{k-1} - \widehat{\mathbf{Y}}_{k-1,r} \right) + \mathbf{Y}_K \left(\mathbf{P}_{\widetilde{\mathbf{Q}}_k} - \mathbf{P}_{\widehat{\mathbf{Q}}_{k,r}} \right). \quad (\text{see (79)}) \end{aligned} \quad (82)$$

Interestingly, the two terms in the last line of (82) are orthogonal, as proved by induction in Appendix F.

Lemma 5. *For every $k \in [2 : K]$, it holds that*

$$\left\langle \mathbf{Y}_{k-1} - \widehat{\mathbf{Y}}_{k-1,r}, \mathbf{Y}_K \left(\mathbf{P}_{\widetilde{\mathbf{Q}}_k} - \mathbf{P}_{\widehat{\mathbf{Q}}_{k,r}} \right) \right\rangle = 0. \quad (83)$$

For fixed $k \in [2 : K]$, Lemma 5 immediately implies that

$$\begin{aligned}
\|\mathbf{Y}_k - \widehat{\mathbf{Y}}_{k,r}\|_F^2 &= \left\| \left(\mathbf{Y}_{k-1} - \widehat{\mathbf{Y}}_{k-1,r} \right) + \mathbf{Y}_K \left(\mathbf{P}_{\widehat{\mathcal{Q}}_k} - \mathbf{P}_{\widehat{\mathcal{Q}}_{k,r}} \right) \right\|_F^2 && \text{(see (82))} \\
&= \|\mathbf{Y}_{k-1} - \widehat{\mathbf{Y}}_{k-1,r}\|_F^2 + \|\mathbf{Y}_K (\mathbf{P}_{\widehat{\mathcal{Q}}_k} - \mathbf{P}_{\widehat{\mathcal{Q}}_{k,r}})\|_F^2 && \text{(see Lemma 5)} \\
&= \|\mathbf{Y}_{k-1} - \widehat{\mathbf{Y}}_{k-1,r}\|_F^2 + \rho_r \left(\widehat{\mathbf{Y}}_{k-1,r} + \mathbf{Y}_k \mathbf{P}_{\mathcal{L}_k} \right). && \text{(see (80) and (78))}
\end{aligned} \tag{84}$$

Recalling from (74) that $\widehat{\mathcal{S}}_{k-1,r} = \text{span}(\widehat{\mathbf{Y}}_{k-1,r})$, we bound the above expression by writing that

$$\begin{aligned}
\|\mathbf{Y}_k - \widehat{\mathbf{Y}}_{k,r}\|_F^2 &= \|\mathbf{Y}_{k-1} - \widehat{\mathbf{Y}}_{k-1,r}\|_F^2 + \rho_r \left(\widehat{\mathbf{Y}}_{k-1,r} + \mathbf{Y}_k \mathbf{P}_{\mathcal{L}_k} \right) \\
&\leq \|\mathbf{Y}_{k-1} - \widehat{\mathbf{Y}}_{k-1,r}\|_F^2 + \left\| \mathbf{P}_{\widehat{\mathcal{S}}_{k-1,r}^\perp} \left(\widehat{\mathbf{Y}}_{k-1,r} + \mathbf{Y}_k \mathbf{P}_{\mathcal{L}_k} \right) \right\|_F^2 \\
&= \|\mathbf{Y}_{k-1} - \widehat{\mathbf{Y}}_{k-1,r}\|_F^2 + \|\mathbf{P}_{\widehat{\mathcal{S}}_{k-1,r}^\perp} \mathbf{y}_k\|_F^2, && \text{(see (74))}
\end{aligned} \tag{85}$$

where the second line follows from the sub-optimality of the choice of subspace $\widehat{\mathcal{S}}_{k-1,r}$. Let us focus on the last norm above. For every k , let $\mathbf{Y}_{k,r} = \text{SVD}_r(\mathbf{Y}_k)$ be a rank- r truncation of \mathbf{Y}_k with the column span $\mathcal{S}_{k,r} = \text{span}(\mathbf{Y}_{k,r})$. We now write that

$$\begin{aligned}
\|\mathbf{P}_{\widehat{\mathcal{S}}_{k-1,r}^\perp} \mathbf{y}_k\|_F &\leq \|\mathbf{P}_{\widehat{\mathcal{S}}_{k-1,r}^\perp} \mathbf{P}_{\mathcal{S}_{k-1,r}} \mathbf{y}_k\|_F + \|\mathbf{P}_{\widehat{\mathcal{S}}_{k-1,r}^\perp} \mathbf{P}_{\mathcal{S}_{k-1,r}^\perp} \mathbf{y}_k\|_F && \text{(triangle inequality)} \\
&\leq \|\mathbf{P}_{\widehat{\mathcal{S}}_{k-1,r}^\perp} \mathbf{P}_{\mathcal{S}_{k-1,r}}\|_F \cdot \|\mathbf{y}_k\| + \|\mathbf{P}_{\mathcal{S}_{k-1,r}^\perp} \mathbf{y}_k\|_F.
\end{aligned} \tag{86}$$

The first norm in the last line above gauges the principal angles between the two r -dimensional subspaces $\widehat{\mathcal{S}}_{k-1,r}$ and $\mathcal{S}_{k-1,r}$. We can bound this norm with a standard perturbation result, for example see [26, Lemma 6] or [55]. More specifically, we may imagine that \mathbf{Y}_{k-1} is a perturbed copy of $\mathbf{Y}_{k-1,r}$. Then the angle between $\mathcal{S}_{k-1,r} = \text{span}(\mathbf{Y}_{k-1,r})$ and $\widehat{\mathcal{S}}_{k-1,r} = \text{span}(\widehat{\mathbf{Y}}_{k-1,r})$ is controlled by the amount of perturbation, namely with the choice of $\mathbf{A} = \widehat{\mathbf{Y}}_{k-1,r}$, $\mathbf{B} = \mathbf{Y}_{k-1}$, $\mathbf{B}_r = \mathbf{Y}_{k-1,r}$ in [26, Lemma 6], we find that

$$\|\mathbf{P}_{\widehat{\mathcal{S}}_{k-1,r}^\perp} \mathbf{P}_{\mathcal{S}_{k-1,r}}\|_F \leq \frac{\|\mathbf{Y}_{k-1} - \widehat{\mathbf{Y}}_{k-1,r}\|_F}{\sigma_r(\mathbf{Y}_{k-1})}. \tag{87}$$

By plugging (87) back into (86), we find that

$$\|\mathbf{P}_{\widehat{\mathcal{S}}_{k-1,r}^\perp} \mathbf{y}_k\| \leq \frac{\|\mathbf{y}_k\|}{\sigma_r(\mathbf{Y}_{k-1})} \cdot \|\mathbf{Y}_{k-1} - \widehat{\mathbf{Y}}_{k-1,r}\|_F + \|\mathbf{P}_{\mathcal{S}_{k-1,r}^\perp} \mathbf{Y}_k\|_F. \tag{88}$$

In turn, for $p > 1$, substituting the above inequality into (85) yields that

$$\begin{aligned}
\|\mathbf{Y}_k - \widehat{\mathbf{Y}}_{k,r}\|_F^2 &\leq \|\mathbf{Y}_{k-1} - \widehat{\mathbf{Y}}_{k-1,r}\|_F^2 + \|\mathbf{P}_{\widehat{\mathcal{S}}_{k-1,r}^\perp} \mathbf{y}_k\|_F^2 && \text{(see (85))} \\
&\leq \left(1 + \frac{p^{\frac{1}{3}} \|\mathbf{y}_k\|^2}{\sigma_r(\mathbf{Y}_{k-1})^2} \right) \|\mathbf{Y}_{k-1} - \widehat{\mathbf{Y}}_{k-1,r}\|_F^2 + \frac{p^{\frac{1}{3}}}{p^{\frac{1}{3}} - 1} \|\mathbf{P}_{\mathcal{S}_{k-1,r}^\perp} \mathbf{y}_k\|_F^2 && \text{(see (88))} \\
&=: \theta_k \|\mathbf{Y}_{k-1} - \widehat{\mathbf{Y}}_{k-1,r}\|_F^2 + \frac{p^{\frac{1}{3}}}{p^{\frac{1}{3}} - 1} \|\mathbf{P}_{\mathcal{S}_{k-1,r}^\perp} \mathbf{y}_k\|_F^2.
\end{aligned} \tag{89}$$

where we used the inequality $(a_1 + a_2)^2 \leq qa_1^2 + \frac{qa_2^2}{q-1}$ for scalars a_1, a_2 and $q > 1$, with the choice of $q = p^{\frac{1}{3}}$. By unfolding the recursion in (89), we arrive at

$$\|\mathbf{Y}_K - \widehat{\mathbf{Y}}_{K,r}\|_F^2 \leq \frac{p^{\frac{1}{3}}}{p^{\frac{1}{3}} - 1} \sum_{k=2}^K \left(\prod_{l=k+1}^K \theta_l \right) \|\mathbf{P}_{\mathcal{S}_{k-1,r}^\perp} \mathbf{y}_k\|_F^2, \tag{90}$$

which completes the proof of Lemma 2.

E Proof of Lemma 4

The proof is by induction. For $k = 1$, it holds that

$$\begin{aligned}
\widehat{\mathbf{Y}}_{1,r} &= \text{SVD}_r(\mathbf{Y}_1) && \text{(see Algorithm 1)} \\
&= \mathbf{Y}_1 \mathbf{P}_{\widehat{\mathcal{Q}}_{1,r}} && \text{(see (74))} \\
&= \mathbf{Y}_K \mathbf{P}_{\mathcal{L}_1} \mathbf{P}_{\widehat{\mathcal{Q}}_{1,r}} \\
&= \mathbf{Y}_K \mathbf{P}_{\widehat{\mathcal{Q}}_{1,r}}, && \left(\widehat{\mathcal{Q}}_{1,r} \subseteq \mathcal{I}_1 \right)
\end{aligned} \tag{91}$$

which proves the base of induction. Next suppose that (77-81) hold for $[2 : k]$ with $k < K$. We now show that (77-81) hold also for $k + 1$. We can then write that

$$\begin{aligned}
\widehat{\mathbf{Y}}_{k+1,r} &= \text{SVD}_r \left(\widehat{\mathbf{Y}}_{k,r} + \begin{bmatrix} \mathbf{0}_{n \times kb} & \mathbf{y}_{k+1} & \mathbf{0}_{n \times (K-k-1)b} \end{bmatrix} \right) && \text{(see Algorithm 1)} \\
&= \text{SVD}_r \left(\mathbf{Y}_K \mathbf{P}_{\widehat{\mathcal{Q}}_{k,r}} + \mathbf{Y}_K \mathbf{P}_{\mathcal{L}_{k+1}} \right) && \text{(assumption of induction)} \\
&= \text{SVD}_r \left(\mathbf{Y}_K \mathbf{P}_{\widehat{\mathcal{Q}}_{k+1}} \right), && \text{(see (79))}
\end{aligned} \tag{92}$$

which completes the proof of Lemma 4.

F Proof of Lemma 5

In this proof only, it is convenient to use the notation $\text{rowspan}(\mathbf{A})$ to denote the row span of a matrix \mathbf{A} , namely $\text{rowspan}(\mathbf{A}) = \text{span}(\mathbf{A}^*)$. For every $k \in [1 : K]$, recall from (80) that $\mathbf{Y}_K(\mathbf{P}_{\widehat{\mathcal{Q}}_k} - \mathbf{P}_{\widehat{\mathcal{Q}}_{k,r}})$ is the residual of rank- r truncation of $\mathbf{Y}_K \mathbf{P}_{\widehat{\mathcal{Q}}_k}$. Consequently,

$$\mathbf{Y}_K(\mathbf{P}_{\widehat{\mathcal{Q}}_k} - \mathbf{P}_{\widehat{\mathcal{Q}}_{k,r}}) = \mathbf{Y}_K \mathbf{P}_{\widehat{\mathcal{Q}}_{k,r}^C}, \quad k \in [1 : K], \tag{93}$$

where $\widehat{\mathcal{Q}}_{k,r}^C$ is the orthogonal complement of $\widehat{\mathcal{Q}}_{k,r}$ with respect to $\widehat{\mathcal{Q}}_k$, namely

$$\widehat{\mathcal{Q}}_k = \widehat{\mathcal{Q}}_{k,r} \oplus \widehat{\mathcal{Q}}_{k,r}^C, \quad \widehat{\mathcal{Q}}_{k,r} \perp \widehat{\mathcal{Q}}_{k,r}^C \quad k \in [1 : K], \tag{94}$$

in which we conveniently set $\widehat{\mathcal{Q}}_1 = \mathcal{I}_1$, see (75). Using (93), we can rewrite (82) as

$$\begin{aligned}
\mathbf{Y}_k - \widehat{\mathbf{Y}}_{k,r} &= (\mathbf{Y}_{k-1} - \widehat{\mathbf{Y}}_{k-1,r}) + \mathbf{Y}_k(\mathbf{P}_{\widehat{\mathcal{Q}}_k} - \mathbf{P}_{\widehat{\mathcal{Q}}_{k,r}}) && \text{(see (82))} \\
&= (\mathbf{Y}_{k-1} - \widehat{\mathbf{Y}}_{k-1,r}) + \mathbf{Y}_K \mathbf{P}_{\widehat{\mathcal{Q}}_{k,r}^C}, && k \in [2 : K].
\end{aligned} \tag{95}$$

With the preliminaries out of the way, let us rewrite the claim of Lemma 5 as

$$\left\langle \mathbf{Y}_{k-1} - \widehat{\mathbf{Y}}_{k-1,r}, \mathbf{Y}_K \mathbf{P}_{\widehat{\mathcal{Q}}_{k,r}^C} \right\rangle = 0, \quad k \in [2 : K], \tag{96}$$

see (83) and (93). Because $\widehat{\mathcal{Q}}_{k,r}^C \subset \widehat{\mathcal{Q}}_k$ by (94), it suffices to instead prove the stronger claim that

$$\text{rowspan}(\mathbf{Y}_{k-1} - \widehat{\mathbf{Y}}_{k-1,r}) \perp \widehat{\mathcal{Q}}_k, \quad k \in [2 : K]. \tag{97}$$

We next prove (97) by induction. The base of induction, namely $k = 2$, is trivial. Suppose now that (97) holds for $[2 : k]$ with $k < K$. We next show that (97) holds for $k + 1$ as well. Note that

$$\begin{aligned} \text{rowspan}(\mathbf{Y}_k - \widehat{\mathbf{Y}}_{k,r}) &= \text{rowspan}\left((\mathbf{Y}_{k-1} - \widehat{\mathbf{Y}}_{k-1,r}) + \mathbf{Y}_K \mathbf{P}_{\widehat{\mathcal{Q}}_{k,r}^C}\right) \quad (\text{see (95)}) \\ &\subseteq \text{rowspan}(\mathbf{Y}_{k-1} - \widehat{\mathbf{Y}}_{k-1,r}) \oplus \widehat{\mathcal{Q}}_{k,r}^C. \end{aligned} \quad (98)$$

As we next show, both subspaces in the last line above are orthogonal to $\widetilde{\mathcal{Q}}_{k+1}$. Indeed, on the one hand,

$$\begin{cases} \text{rowspan}(\mathbf{Y}_{k-1} - \widehat{\mathbf{Y}}_{k-1,r}) \perp \widetilde{\mathcal{Q}}_k \supseteq \widehat{\mathcal{Q}}_{k,r}, & (\text{induction hypothesis and (81)}) \\ \text{rowspan}(\mathbf{Y}_{k-1} - \widehat{\mathbf{Y}}_{k-1,r}) \subset \mathcal{J}_{k-1} \perp \mathcal{I}_{k+1}, & (\text{see (81) and (76)}) \end{cases} \\ \implies \text{rowspan}(\mathbf{Y}_{k-1} - \widehat{\mathbf{Y}}_{k-1,r}) \perp (\widehat{\mathcal{Q}}_{k,r} \oplus \mathcal{I}_{k+1}) = \widetilde{\mathcal{Q}}_{k+1}. \quad (\text{see (79)}) \quad (99)$$

On the other hand,

$$\begin{cases} \widehat{\mathcal{Q}}_{k,r}^C \perp \widehat{\mathcal{Q}}_{k,r}, \\ \widehat{\mathcal{Q}}_{k,r}^C \subset \widetilde{\mathcal{Q}}_k \subset \mathcal{J}_k \perp \mathcal{I}_{k+1}, & (\text{see (81) and (76)}) \end{cases} \\ \implies \widehat{\mathcal{Q}}_{k,r}^C \perp (\widehat{\mathcal{Q}}_{k,r} \oplus \mathcal{I}_{k+1}) = \widetilde{\mathcal{Q}}_{k+1}. \quad (\text{see (79)}) \quad (100)$$

By combining (99) and (100), we conclude that

$$\begin{aligned} \text{rowspan}(\mathbf{Y}_k - \widehat{\mathbf{Y}}_{k,r}) &\subseteq \text{rowspan}(\mathbf{Y}_{k-1} - \widehat{\mathbf{Y}}_{k-1,r}) \oplus \widehat{\mathcal{Q}}_{k,r}^C \quad (\text{see (98)}) \\ &\perp \widetilde{\mathcal{Q}}_{k+1}. \quad (\text{see (99,100)}) \end{aligned} \quad (101)$$

Therefore, (97) holds for every $k \in [2 : K]$ by induction. In particular, this proves Lemma 5.

G Proof of Lemma 3

Recall that $\mathbf{y}_k \in \mathbb{R}^{n \times b}$, $\mathbf{Y}_k \in \mathbb{R}^{n \times kb}$ denote the k th block and the concatenation of the first k blocks of data, respectively. Since the data vectors are independently drawn from a zero-mean Gaussian probability measure with covariance matrix Ξ , it follows from (55,56) that

$$\begin{aligned} \mathbf{y}_k &= \mathbf{S}\Sigma\mathbf{g}_k, \\ \mathbf{Y}_k &= \mathbf{S}\Sigma\mathbf{G}_k, \end{aligned} \quad (102)$$

for every $k \in [1 : K]$, where $\mathbf{g}_k \in \mathbb{R}^{n \times b}$ and $\mathbf{G}_k \in \mathbb{R}^{n \times kb}$ are standard random Gaussian matrices. For fixed $k \in [2 : K]$, let us now study each of the random quantities on the right-hand side of (69). The following results are proved in Appendices H and I, respectively.

Lemma 6. (Bound on $\|\mathbf{y}_k\|$) For $\alpha \geq 1$, $p > 1$, and fixed $k \in [1 : K]$, it holds that

$$\|\mathbf{y}_k\| \leq p^{\frac{1}{6}}(\sigma_1 + \sqrt{2}\alpha p^{-\frac{1}{6}}\rho_r)\sqrt{b}, \quad (103)$$

except with a probability of at most $e^{-C\alpha^2 b}$ and provided that

$$b \geq \frac{\alpha^2 r}{(p^{\frac{1}{6}} - 1)^2}. \quad (104)$$

Lemma 7. (Bound on $\sigma_r(\mathbf{Y}_k)$) For $\alpha \geq 1$, $p > 1$, and fixed $k \in [1 : K]$, it holds that

$$\sigma_r(\mathbf{Y}_k) \geq p^{-\frac{1}{6}}\sigma_r\sqrt{kb}, \quad (105)$$

except with a probability of at most $e^{-C\alpha^2 r}$ and provided that

$$b \geq \frac{\alpha^2 r}{(1 - p^{-\frac{1}{6}})^2}. \quad (106)$$

By combining Lemmas 6 and 7, we find for fixed $k \in [2 : K]$ that

$$\begin{aligned} \theta_k &= 1 + \frac{p^{\frac{1}{3}} \|\mathbf{y}_k\|^2}{\sigma_r(\mathbf{Y}_{k-1})^2} \quad (\text{see (68)}) \\ &\leq 1 + \frac{p(\sigma_1 + \sqrt{2}\alpha p^{-\frac{1}{6}}\rho_r)^2 b}{\sigma_r^2(k-1)b} \quad (\text{see Lemmas 6 and 7}) \\ &=: 1 + \frac{p\eta_r^2}{k-1}, \end{aligned} \quad (107)$$

except with a probability of at most $e^{-C\alpha^2 r}$ and provided that (106) holds. In particular, it follows that

$$\begin{aligned} \prod_{l=k+1}^K \theta_l &\leq \prod_{l=k+1}^K \left(1 + \frac{p\eta_r^2}{l-1}\right) \quad (\text{see (107)}) \\ &\leq \frac{(K-1 + p\eta_r^2)^{K-1+p\eta_r^2}}{(K-1)^{K-1}} \cdot \frac{(k-1)^{k-1}}{(k-1 + p\eta_r^2)^{k-1+p\eta_r^2}} \quad (\text{see below}) \\ &= \left(1 + \frac{p\eta_r^2}{K-1}\right)^{K-1} \left(1 + \frac{p\eta_r^2}{k-1}\right)^{-k+1} \left(\frac{K-1 + p\eta_r^2}{k-1 + p\eta_r^2}\right)^{p\eta_r^2}, \end{aligned} \quad (108)$$

holds for every $k \in [2 : K]$ and except with a probability of at most $Ke^{-C\alpha r}$, where the failure probability follows from an application of the union bound. The second line above is obtained by bounding the logarithm of the product in that line with the corresponding integral. More specifically, it holds that

$$\begin{aligned} &\log \left(\prod_{l=k+1}^K \left(1 + \frac{p\eta_r^2}{l-1}\right) \right) \\ &= \sum_{l=k}^{K-1} \log \left(1 + \frac{p\eta_r^2}{l}\right) \\ &\leq \int_{k-1}^{K-1} \log \left(1 + \frac{p\eta_r^2}{x}\right) dx \\ &= (K-1 + p\eta_r^2) \log(K-1 + p\eta_r^2) - (K-1) \log(K-1) \\ &\quad - (k-1 + p\eta_r^2) \log(k-1 + p\eta_r^2) + (k-1) \log(k-1), \end{aligned} \quad (109)$$

where the third line above follows because the integrand is decreasing in x . Let us further simplify (108). Note that $K \geq k \geq 2$ and that $p\eta_r^2 \geq 1$ by its definition in (107). Consequently, using the relation $2 \leq (1 + 1/x)^x \leq e$ for $x \geq 1$, we can write that

$$2 \leq \left(1 + \frac{p\eta_r^2}{k-1}\right)^{\frac{k-1}{p\eta_r^2}} \leq e, \quad 2 \leq \left(1 + \frac{p\eta_r^2}{K-1}\right)^{\frac{K-1}{p\eta_r^2}} \leq e. \quad (110)$$

In turn, (110) allows us to simplify (108) as follows:

$$\begin{aligned} \prod_{l=k+1}^K \theta_l &\leq \left(1 + \frac{p\eta_r^2}{K-1}\right)^{K-1} \left(1 + \frac{p\eta_r^2}{k-1}\right)^{-k+1} \left(\frac{K-1+p\eta_r^2\eta_r}{k-1+p\eta_r^2}\right)^{p\eta_r^2} \quad (\text{see (108)}) \\ &\leq \left(\frac{e}{2}\right)^{p\eta_r^2} \left(\frac{K-1+p\eta_r^2}{k-1+p\eta_r^2}\right)^{p\eta_r^2}. \quad (\text{see (110)}) \end{aligned} \quad (111)$$

Next we control the random variable $\|\mathbf{P}_{S_{k-1}^\perp} \mathbf{y}_k\|_F$ in (69) with the following result, proved in Appendix J.

Lemma 8. (Bound on the Innovation) *For $\alpha \geq 1$ and fixed $k \in [2 : K]$, it holds that*

$$\|\mathbf{P}_{S_{k-1}^\perp, r} \mathbf{y}_k\|_F \leq 5\alpha \min(\kappa_r \rho_r, \sqrt{r}\sigma_1 + \rho_r) \sqrt{b}, \quad (112)$$

except with a probability of at most $e^{-C\alpha^2 r}$ and provided that $b \geq C\alpha^2 r$.

By combining Lemma 8 and (111), we finally find a stochastic bound for the right-hand side of (69). More specifically, it holds that

$$\begin{aligned} &\|\mathbf{Y}_K - \widehat{\mathbf{Y}}_{K,r}\|_F^2 \\ &\leq \frac{p^{\frac{1}{3}}}{p^{\frac{1}{3}} - 1} \sum_{k=2}^K \left(\prod_{l=k+1}^K \theta_l \right) \|\mathbf{P}_{S_{k-1}^\perp, r} \mathbf{y}_k\|_F^2 \quad (\text{see (69)}) \\ &\leq \frac{50p^{\frac{1}{3}}\alpha^2}{p^{\frac{1}{3}} - 1} \min(\kappa_r^2 \rho_r^2, r\sigma_1^2 + \rho_r^2) b \cdot \left(\frac{e}{2}\right)^{p\eta_r^2} (K-1+p\eta_r^2)^{p\eta_r^2} \sum_{k=2}^K (k-1+p\eta_r^2)^{-p\eta_r^2} \quad (\text{see (111) and Lemma 8}) \\ &\leq \frac{50p^{\frac{1}{3}}\alpha^2}{p^{\frac{1}{3}} - 1} \min(\kappa_r^2 \rho_r^2, r\sigma_1^2 + \rho_r^2) b \cdot \left(\frac{e}{2}\right)^{p\eta_r^2} (K-1+p\eta_r^2)^{p\eta_r^2} \int_{p\eta_r^2}^{\infty} x^{-p\eta_r^2} dx \\ &= \frac{50p^{\frac{1}{3}}\alpha^2}{p^{\frac{1}{3}} - 1} \min(\kappa_r^2 \rho_r^2, r\sigma_1^2 + \rho_r^2) b \cdot \left(\frac{e}{2}\right)^{p\eta_r^2} (K-1+p\eta_r^2)^{p\eta_r^2} \cdot \frac{(p\eta_r^2)^{-p\eta_r^2+1}}{p\eta_r^2 - 1} \\ &\leq \frac{50p^{\frac{1}{3}}\alpha^2}{p^{\frac{1}{3}} - 1} \min(\kappa_r^2 \rho_r^2, r\sigma_1^2 + \rho_r^2) b \left(\frac{2K}{p\eta_r^2} + 2\right)^{p\eta_r^2} \frac{p\eta_r^2}{p\eta_r^2 - 1} \\ &\leq \frac{50p^{\frac{4}{3}}\alpha^2}{(p^{\frac{1}{3}} - 1)^2} \cdot \min(\kappa_r^2 \rho_r^2, r\sigma_1^2 + \rho_r^2) \eta_r^2 b \left(\frac{2K}{p\eta_r^2} + 2\right)^{p\eta_r^2}, \quad (p, \eta_r \geq 1) \end{aligned} \quad (113)$$

except with a probability of at most $e^{-C\alpha^2 r}$ and provided that

$$b \geq \frac{p^{\frac{1}{3}}\alpha^2 r}{(p^{\frac{1}{6}} - 1)^2}, \quad b \geq C\alpha^2 r.$$

This completes the proof of Lemma 3.

H Proof of Lemma 6

Note that

$$\begin{aligned} \|\mathbf{y}_k\| &= \|\mathbf{S}\boldsymbol{\Sigma}\mathbf{g}_k\| \quad (\text{see (102)}) \\ &= \|\boldsymbol{\Sigma}\mathbf{g}_k\| \quad (\mathbf{S}^*\mathbf{S} = \mathbf{I}_n) \\ &\leq \|\boldsymbol{\Sigma}[1 : r, 1 : r] \cdot \mathbf{g}_k[1 : r, :]\| + \|\boldsymbol{\Sigma}[r+1 : n, r+1 : n] \cdot \mathbf{g}_k[r+1 : n, :]\| \quad (\text{triangle inequality}) \\ &\leq \sigma_1 \cdot \|\mathbf{g}_k[1 : r, :]\| + \|\boldsymbol{\Sigma}[r+1 : n, r+1 : n] \cdot \mathbf{g}_k[r+1 : n, :]\| \\ &\leq \sigma_1 \cdot \|\mathbf{g}_k[1 : r, :]\| + \|\boldsymbol{\Sigma}[r+1 : n, r+1 : n] \cdot \mathbf{g}_k[r+1 : n, :]\|_F, \end{aligned} \quad (114)$$

where we used MATLAB's matrix notation as usual. Note that both $\mathbf{g}_k[1:r, :] \in \mathbb{R}^{r \times b}$ and $\mathbf{g}_k[r+1:n, :] \in \mathbb{R}^{(n-r) \times b}$ in (114) are standard Gaussian random matrices. For $\alpha \geq 1$ and $p > 1$, invoking the results about the spectrum of Gaussian random matrices in Appendix A yields that

$$\begin{aligned}
\|\mathbf{y}_k\| &\leq \sigma_1 \cdot \|\mathbf{g}_k[1:r, :]\| + \|\boldsymbol{\Sigma}[r+1:n, r+1:n] \cdot \mathbf{g}_k[r+1:n, :]\|_F \quad (\text{see (114)}) \\
&\leq \sigma_1(\sqrt{b} + \alpha\sqrt{r}) + \sqrt{2}\alpha\|\boldsymbol{\Sigma}[r+1:n, r+1:n]\|_F\sqrt{b} \quad (\text{see (51,53) and } b \geq r) \\
&= \sigma_1(\sqrt{b} + \alpha\sqrt{r}) + \alpha\rho_r\sqrt{2b} \quad (\text{see (56,57)}) \\
&\leq p^{\frac{1}{6}}\sigma_1\sqrt{b} + \alpha\rho_r\sqrt{2b}, \quad \left(\text{if } b \geq \frac{\alpha^2 r}{(p^{\frac{1}{6}} - 1)^2}\right) \tag{115}
\end{aligned}$$

except with a probability of at most $e^{-C\alpha^2 r} + e^{-C\alpha^2 b} \leq e^{-C\alpha^2 r}$, where this final inequality follows from the assumption that $b \geq r$. This completes the proof of Lemma 6. We remark that a slightly stronger bound can be obtained by using Slepian's inequality for comparing Gaussian processes, see [53, Section 5.3.1] and [56, Section 3.1].

I Proof of Lemma 7

For a matrix $\mathbf{A} \in \mathbb{R}^{n \times kb}$, it follows from the Fisher-Courant representation of the singular values that

$$\sigma_r(\mathbf{A}) \geq \sigma_r(\mathbf{A}[1:r, :]). \tag{116}$$

Alternatively, (116) might be verified using Cauchy's interlacing theorem applied to $\mathbf{A}\mathbf{A}^*$. For a vector $\gamma \in \mathbb{R}^{r \times r}$ and matrix $\mathbf{A} \in \mathbb{R}^{r \times r}$, we also have the useful inequality

$$\sigma_r(\text{diag}(\gamma)\mathbf{A}) \geq \min_{i \in [r]} |\gamma[i]| \cdot \sigma_r(\mathbf{A}), \tag{117}$$

where $\text{diag}(\gamma) \in \mathbb{R}^{r \times r}$ is the diagonal matrix formed from the entries of γ . Using the above inequalities, we may write that

$$\begin{aligned}
\sigma_r(\mathbf{Y}_k) &= \sigma_r(\mathbf{S}\boldsymbol{\Sigma}\mathbf{G}_k) \quad (\text{see (102)}) \\
&= \sigma_r(\boldsymbol{\Sigma}\mathbf{G}_k) \quad (\mathbf{S}^*\mathbf{S} = \mathbf{I}_n) \\
&\geq \sigma_r(\boldsymbol{\Sigma}[1:r, 1:r] \cdot \mathbf{G}_k[1:r, :]) \quad (\text{see (116)}) \\
&\geq \sigma_r \cdot \sigma_r(\mathbf{G}_k[1:r, :]). \quad (\text{see (117,56)}) \tag{118}
\end{aligned}$$

Note also that $\mathbf{G}_k[1:r, :] \in \mathbb{R}^{r \times kb}$ above is a standard Gaussian random matrix. Using the spectral properties listed in Appendix A, we can therefore write that

$$\begin{aligned}
\sigma_r(\mathbf{Y}_k) &\geq \sigma_r \cdot \sigma_r(\mathbf{G}_k[1:r, :]) \quad (\text{see (118)}) \\
&\geq \sigma_r \cdot (\sqrt{kb} - \alpha\sqrt{r}) \quad (\text{see (51) and } b \geq r) \\
&\geq \sigma_r \cdot p^{-\frac{1}{6}}\sqrt{kb}, \quad \left(\text{if } b \geq \frac{\alpha^2 r}{(1 - p^{-\frac{1}{6}})^2}\right) \tag{119}
\end{aligned}$$

except with a probability of at most $e^{-C\alpha^2 r}$. This completes the proof of Lemma 7.

J Proof of Lemma 8

Without loss of generality, we set $\mathbf{S} = \mathbf{I}_n$ in (55) to simplify the presentation, as this renders the contribution of the bottom rows of \mathbf{y}_k to the innovation typically small. We first separate this term via the inequality

$$\begin{aligned} \|\mathbf{P}_{\mathcal{S}_{k-1,r}^\perp} \mathbf{y}_k\|_F &= \left\| \mathbf{P}_{\mathcal{S}_{k-1,r}^\perp} \begin{bmatrix} \mathbf{y}_k[1:r,:] \\ \mathbf{y}_k[r+1:n,:] \end{bmatrix} \right\|_F \\ &\leq \left\| \mathbf{P}_{\mathcal{S}_{k-1,r}^\perp} \begin{bmatrix} \mathbf{y}_k[1:r,:] \\ \mathbf{0}_{(n-r) \times b} \end{bmatrix} \right\|_F + \|\mathbf{y}_k[r+1:n,:]\|_F. \quad (\text{triangle inequality}) \end{aligned} \quad (120)$$

To control the last norm above, we simply write that

$$\begin{aligned} \|\mathbf{y}_k[r+1:n,:]\|_F &= \|\boldsymbol{\Sigma}[r+1:n, r+1:n] \cdot \mathbf{g}_k[r+1:n,:]\|_F \quad (\text{see (102)}) \\ &\leq \alpha \|\boldsymbol{\Sigma}[r+1:n, r+1:n]\|_F \sqrt{2b} \quad (\text{see (53)}) \\ &= \alpha \rho_r \sqrt{2b}, \quad (\text{see (57)}) \end{aligned} \quad (121)$$

except with a probability of at most $e^{-C\alpha^2 b}$. In the second line above, we used the fact that \mathbf{g}_k is a standard Gaussian random matrix. It therefore remains to control the first norm in the last line of (120). Note that

$$\begin{aligned} \left\| \mathbf{P}_{\mathcal{S}_{k-1,r}^\perp} \begin{bmatrix} \mathbf{y}_k[1:r,:] \\ \mathbf{0}_{(n-r) \times b} \end{bmatrix} \right\|_F &= \left\| \mathbf{P}_{\mathcal{S}_{k-1,r}^\perp} \begin{bmatrix} \mathbf{I}_r & \\ & \mathbf{0}_{n-r} \end{bmatrix} \cdot \begin{bmatrix} \mathbf{y}_k[1:r,:] \\ \mathbf{0}_{(n-r) \times b} \end{bmatrix} \right\|_F \\ &=: \left\| \mathbf{P}_{\mathcal{S}_{k-1,r}^\perp} \mathbf{J}_r \cdot \begin{bmatrix} \mathbf{y}_k[1:r,:] \\ \mathbf{0}_{(n-r) \times b} \end{bmatrix} \right\|_F \\ &\leq \|\mathbf{P}_{\mathcal{S}_{k-1,r}^\perp} \mathbf{J}_r\|_F \cdot \|\mathbf{y}_k[1:r,:]\| \\ &\leq \|\mathbf{P}_{\mathcal{S}_{k-1,r}^\perp} \mathbf{J}_r\|_F \cdot \|\boldsymbol{\Sigma}[1:r, 1:r]\| \cdot \|\mathbf{g}_k[1:r,:]\| \quad (\text{see (102)}) \\ &\leq \|\mathbf{P}_{\mathcal{S}_{k-1,r}^\perp} \mathbf{J}_r\|_F \cdot \sigma_1 \cdot (\sqrt{b} + \alpha\sqrt{r}) \quad (\text{see (56,51)}) \\ &\leq \|\mathbf{P}_{\mathcal{S}_{k-1,r}^\perp} \mathbf{J}_r\|_F \cdot \sigma_1 \sqrt{2b}, \quad (\text{if } b \geq C\alpha^2 r) \end{aligned} \quad (122)$$

except with a probability of at most $e^{-C\alpha^2 r}$ and provided that $b \geq C\alpha^2 r$. The fifth line above again uses the fact that \mathbf{g}_k is a standard Gaussian random matrix. Let us now estimate the norm in the last line above. Recall that $\mathbf{P}_{\mathcal{S}_{k-1,r}^\perp} \in \mathbb{R}^{n \times n}$ projects onto the span of $\mathbf{Y}_{k-1,r} = \text{SVD}_r(\mathbf{Y}_{k-1})$, namely $\mathbf{P}_{\mathcal{S}_{k-1,r}^\perp}$ projects onto the span of leading r left singular vectors of $\mathbf{Y}_{k-1} = \boldsymbol{\Sigma} \mathbf{G}_{k-1}$, see (102). Because the diagonal entries of $\boldsymbol{\Sigma} \in \mathbb{R}^{n \times n}$ are in nonincreasing order, it is natural to expect that $\mathbf{P}_{\mathcal{S}_{k-1,r}^\perp} \approx \mathbf{J}_r$. We now formalise this notion using standard results from the perturbation theory. Note that one might think of $\mathbf{Y}_{k-1,r} = \text{SVD}_r(\mathbf{Y}_{k-1})$ as a perturbed copy of \mathbf{Y}_{k-1} . Note also that \mathbf{J}_r is the orthogonal projection onto the subspace

$$\text{span} \left(\begin{bmatrix} \mathbf{Y}_{k-1}[1:r,:] \\ \mathbf{0}_{(n-r) \times (k-1)b} \end{bmatrix} \right),$$

because $\mathbf{Y}_{k-1}[1:r, :]$ is almost surely full-rank. An application of Lemma 6 in [26] with \mathbf{A} as specified inside the parenthesis above and $\mathbf{B} = \mathbf{Y}_{k-1}$ yields that

$$\begin{aligned}
\|\mathbf{P}_{\mathcal{S}_{k-1,r}^\perp} \mathbf{J}_r\|_F &\leq \frac{\left\| \mathbf{Y}_{k-1} - \begin{bmatrix} \mathbf{Y}_{k-1}[1:r, :] \\ \mathbf{0}_{(n-r) \times (k-1)b} \end{bmatrix} \right\|_F}{\sigma_r(\mathbf{Y}_{k-1})} \\
&= \frac{\|\mathbf{Y}_{k-1}[r+1:n, :]\|_F}{\sigma_r(\mathbf{Y}_{k-1})} \\
&= \frac{\|\boldsymbol{\Sigma}[r+1:n, r+1:n] \cdot \mathbf{G}_{k-1}[r+1:n, :]\|_F}{\sigma_r(\mathbf{Y}_{k-1})} \quad (\text{see (102)}) \\
&\leq \frac{\alpha \|\boldsymbol{\Sigma}[r+1:n, r+1:n]\|_F \sqrt{2(k-1)b}}{\sigma_r \sqrt{(k-1)b/2}} \quad (\text{see (53) and Lemma 7 with } p=8) \\
&= \frac{2\alpha\rho_r}{\sigma_r}, \quad (\text{see (57)}) \tag{123}
\end{aligned}$$

provided that $b \geq C\alpha^2 r$ and except with a probability of at most $e^{-C\alpha^2 b} + e^{-C\alpha^2 r} \leq e^{-C\alpha^2 r}$, where this last inequality follows from the assumption that $b \geq r$. It also trivially holds that

$$\|\mathbf{P}_{\mathcal{S}_{k-1,r}^\perp} \mathbf{J}_r\|_F \leq \|\mathbf{P}_{\mathcal{S}_{k-1,r}^\perp}\| \cdot \|\mathbf{J}_r\|_F \leq \|\mathbf{J}_r\|_F = \|\mathbf{I}_r\|_F = \sqrt{r},$$

where we used above the definition of \mathbf{J}_r in (122). Therefore, overall we find that

$$\|\mathbf{P}_{\mathcal{S}_{k-1,r}^\perp} \mathbf{J}_r\|_F \leq \min\left(\frac{2\alpha\rho_r}{\sigma_r}, \sqrt{r}\right). \tag{124}$$

Substituting the above bound back into (122) yields that

$$\begin{aligned}
\left\| \mathbf{P}_{\mathcal{S}_{k-1,r}^\perp} \begin{bmatrix} \mathbf{y}_k[1:r, :] \\ \mathbf{0}_{(n-r) \times b} \end{bmatrix} \right\|_F &\leq \|\mathbf{P}_{\mathcal{S}_{k-1,r}^\perp} \mathbf{J}_r\|_F \cdot \sigma_1 \sqrt{2b} \quad (\text{see (122)}) \\
&\leq \min(\alpha\kappa_r \rho_r, \sigma_1 \sqrt{r}) \sqrt{8b}, \quad (\text{see (124,57)}) \tag{125}
\end{aligned}$$

except with a probability of at most $e^{-C\alpha^2 r}$. Combining (121) and (125) finally controls the innovation as

$$\begin{aligned}
\|\mathbf{P}_{\mathcal{S}_{k-1,r}^\perp} \mathbf{y}_k\|_F &\leq \left\| \mathbf{P}_{\mathcal{S}_{k-1,r}^\perp} \begin{bmatrix} \mathbf{y}_k[1:r, :] \\ \mathbf{0}_{(n-r) \times b} \end{bmatrix} \right\|_F + \|\mathbf{y}_k[r+1:n, :]\|_F \quad (\text{see (120)}) \\
&\leq \min(\alpha\kappa_r \rho_r, \sigma_1 \sqrt{r}) \sqrt{8b} + \alpha\rho_r \sqrt{2b} \quad (\text{see (125,121)}) \\
&\leq 5\alpha \min(\kappa_r \rho_r, \sigma_1 \sqrt{r} + \rho_r) \sqrt{b}, \quad (\alpha, \kappa_r \geq 1) \tag{126}
\end{aligned}$$

except with a probability of at most $e^{-C\alpha^2 r}$ and provided that $b \geq C\alpha^2 r$. This completes the proof of Lemma 8.