

Streaming PCA: Matching Matrix Bernstein and Near-Optimal Finite Sample Guarantees for Oja’s Algorithm

Prateek Jain* Chi Jin† Sham M. Kakade‡ Praneeth Netrapalli§ Aaron Sidford¶

March 29, 2016

Abstract

This work provides improved guarantees for streaming principle component analysis (PCA). Given $\mathbf{A}_1, \dots, \mathbf{A}_n \in \mathbb{R}^{d \times d}$ sampled independently from distributions satisfying $\mathbb{E}[\mathbf{A}_i] = \Sigma$ for $\Sigma \succeq \mathbf{0}$, this work provides an $O(d)$ -space linear-time single-pass streaming algorithm for estimating the top eigenvector of Σ . The algorithm nearly matches (and in certain cases improves upon) the accuracy obtained by the standard batch method that computes top eigenvector of the empirical covariance $\frac{1}{n} \sum_{i \in [n]} \mathbf{A}_i$ as analyzed by the matrix Bernstein inequality. Moreover, to achieve constant accuracy, our algorithm improves upon the best previous known sample complexities of streaming algorithms by either a multiplicative factor of $O(d)$ or $1/\text{gap}$ where gap is the relative distance between the top two eigenvalues of Σ .

These results are achieved through a novel analysis of the classic Oja’s algorithm, one of the oldest and most popular algorithms for streaming PCA. In particular, this work shows that simply picking a random initial point \mathbf{w}_0 and applying the update rule $\mathbf{w}_{i+1} = \mathbf{w}_i + \eta_i \mathbf{A}_i \mathbf{w}_i$ suffices to accurately estimate the top eigenvector, with a suitable choice of η_i . We believe our result sheds light on how to efficiently perform streaming PCA both in theory and in practice, and we hope that our analysis may serve as the basis for analyzing many variants and extensions of streaming PCA.

*Microsoft Research India. Email: prajain@microsoft.com

†UC Berkeley. Email: chijin@cs.berkeley.edu

‡University of Washington. Email: sham@cs.washington.edu

§Microsoft Research New England. Email: praneeth@microsoft.com

¶Microsoft Research New England. Email: asid@microsoft.com

1 Introduction

Principal component analysis (PCA) is one of the most fundamental problems in machine learning, numerical linear algebra, and data analysis. It is commonly used for data compression, image processing, and visualization [1] etc.

When we desire to perform PCA on large data sets, it may be the case that we cannot afford more than single pass over the data (or worse to even store the data in the first place) [2, 3, 4]. To alleviate this issue, a popular line of research over the past several decades has been to consider streaming algorithms for PCA under the assumption that the data has reasonable statistical properties [5, 6, 7, 8, 9]. There have been significant breakthroughs in getting near-optimal streaming PCA algorithms under fairly specialized models, e.g. spiked covariance [9].

This work considers one of the most natural variants of PCA, estimating the top eigenvector of a symmetric matrix, under a mild (and standard) set of assumptions under which concentration of measure applies (under the matrix Bernstein inequality [10, 11]). In particular, the setting is as follows:

Definition 1 (Streaming PCA). Let $\mathbf{A}_1, \mathbf{A}_2, \dots, \mathbf{A}_n \in \mathbb{R}^{d \times d}$ be a sequence of (not necessarily symmetric) matrices sampled independently from distributions that satisfy the following:

1. $\mathbb{E}[\mathbf{A}_i] = \Sigma$ for symmetric positive semidefinite (PSD) matrix $\Sigma \in \mathbb{R}^{d \times d}$,
2. $\|\mathbf{A}_i - \Sigma\|_2 \leq \mathcal{M}$ with probability 1, and
3. $\max \left\{ \left\| \mathbb{E} \left[(\mathbf{A}_i - \Sigma)(\mathbf{A}_i - \Sigma)^\top \right] \right\|_2, \left\| \mathbb{E} \left[(\mathbf{A}_i - \Sigma)^\top (\mathbf{A}_i - \Sigma) \right] \right\|_2 \right\} \leq \mathcal{V}$.

Let $\mathbf{v}_1, \dots, \mathbf{v}_d$ denote the eigenvectors of Σ and $\lambda_1 \geq \dots \geq \lambda_d$ denote the corresponding eigenvalues. Our goal is to compute an ϵ -approximation to \mathbf{v}_1 , that is a unit vector \mathbf{w} such that $\sin^2(\mathbf{w}, \mathbf{v}_1) \triangleq 1 - (\mathbf{w}^\top \mathbf{v}_1)^2 \leq \epsilon$, in a single pass while minimizing space, time, and error (i.e. ϵ). Note that $\sin(\mathbf{w}, \mathbf{v}_1)$ denotes the sin of the angle between \mathbf{w} and \mathbf{v}_1 .

A special case of Streaming PCA is to estimate the top eigenvector of the covariance matrix of a distribution \mathcal{D} over \mathbb{R}^d , i.e. given independent samples $\mathbf{a}_1, \dots, \mathbf{a}_n \in \mathbb{R}^d$ estimate the top eigenvector of $\mathbb{E}_{\mathbf{a} \sim \mathcal{D}}[\mathbf{a}\mathbf{a}^\top]$. This encompasses the popular "spiked covariance model" [12].

It is well known that to solve the Streaming PCA problem, one can simply compute the empirical covariance matrix $\frac{1}{n} \sum_{i \in [n]} \mathbf{A}_i$ and compute the right singular vector of this matrix. Here, matrix Bernstein inequality [10, 11] and Wedin's theorem [13] implies the following standard sample complexity bound for the Streaming PCA problem:

Theorem 1.1 (Eigenvector Concentration using matrix Bernstein and Wedin's theorem). *Under the assumptions of Definition 1, the top right singular vector $\hat{\mathbf{v}}$ of $\hat{\Sigma} = \frac{1}{n} \sum_{i \in [n]} \mathbf{A}_i$ is an ϵ -approximation to the top eigenvector \mathbf{v}_1 of Σ with probability $1 - \delta$, where*

$$\sin^2(\hat{\mathbf{v}}, \mathbf{v}_1) \leq \epsilon \leq \frac{16\mathcal{V} \log \frac{d}{\delta}}{(\lambda_1 - \lambda_2)^2} \cdot \frac{1}{n} + \left(\frac{4\mathcal{M} \log \frac{d}{\delta}}{\lambda_1 - \lambda_2} \right)^2 \cdot \frac{1}{n^2}.$$

Theorem 1.1 is essentially the previous best sample complexity known for estimating the top eigenvector¹. Unfortunately, the above is purely a statistical claim, and, algorithmically, there are least two concerns.

¹In recent work in [14] it was shown that the $\log(d/\delta)$ factor in the first term could be removed asymptotically for small enough ϵ if only constant success probability is required.

First, computing the empirical covariance matrix $\widehat{\Sigma} = \frac{1}{n} \sum_{i \in [n]} \mathbf{A}_i$ naively requires $O(d^2)$ time and space, and second, computing the top eigenvector of the empirical covariance matrix in general may require super linear time [15]. While there have been many attempts to produce streaming algorithms that use only $O(d)$ space to solve the streaming PCA problem, to our knowledge, all previous methods either lose a multiplicative factor of either $\frac{\lambda_1}{\lambda_1 - \lambda_2}$ or d in the analysis in order to achieve constant accuracy when applied in our setting [7, 8, 16, 9, 14].

In an attempt to overcome this limitation and improve the guarantees for solving the streaming PCA problem, this work seeks to address the following question:

Can we match the sample complexity of matrix Bernstein + Wedin’s theorem with an algorithm that uses $O(d)$ space only and takes a single linear-time pass over the input?

This work answers this question in the affirmative, showing that one can succeed with constant probability matching the sample complexity of Theorem 1.1 up to logarithmic terms and small additive factors. Interestingly, this is achieved by providing a novel analysis of the classical Oja’s algorithm, which is perhaps, the most popular algorithm for Streaming PCA [6].

Algorithm 1 Oja’s algorithm for computing top eigenvector

Input: $\mathbf{A}_1, \dots, \mathbf{A}_n$.

Choose \mathbf{w}_0 uniformly at random from the unit sphere

for $t = 1, \dots, n$ **do**

$\mathbf{w}_i \leftarrow \mathbf{w}_{i-1} + \eta_i \mathbf{A}_i \mathbf{w}_{i-1}$

$\mathbf{w}_i \leftarrow \mathbf{w}_i / \|\mathbf{w}_i\|_2$

end for

Output: \mathbf{w}_n

Oja’s algorithm is one of the simplest algorithms one would imagine for the streaming PCA problem (See Algorithm 1). In fact, due to its simplicity, it was proposed a neurally plausible algorithm. In the case that each \mathbf{A}_i comes from the same distribution \mathcal{D} it corresponds to simply performing projected stochastic gradient descent on the objective function of maximizing the Rayleigh Quotient over the distribution $\max_{\|w\|_2=1} \mathbb{E}_{\mathbf{A} \sim \mathcal{D}} \mathbf{w}^\top \mathbf{A} \mathbf{w}$. It is well known that under very mild conditions on the stepsize sequence, Oja’s algorithm asymptotically converges to the top eigenvector of the covariance matrix Σ [6]. However, obtaining optimal rates of convergence, let alone finite sample guarantees, for Streaming PCA has been quite challenging. The best known results are off from Theorem 1.1 by a factor of $\mathcal{O}(d)$ [9].

This work shows that for proper choice of learning rates η_i , Oja’s algorithm in fact can improve the best known results for streaming PCA and answer our question in the affirmative. In particular, we have that:

Theorem 1.2. *Let the assumptions of Definition 1 hold. Suppose the step size sequence for Algorithm 1 is chosen to be $\eta_i = \frac{\log d}{(\lambda_1 - \lambda_2)(\beta + i)}$, where*

$$\beta \triangleq 40 \max \left(\frac{\mathcal{M} \log d}{(\lambda_1 - \lambda_2)}, \frac{(\mathcal{V} + (\lambda_1)^2) \log^2 d}{(\lambda_1 - \lambda_2)^2} \right).$$

Then the output \mathbf{w}_n of Algorithm 1 is an ϵ -approximation to the top eigenvector \mathbf{v}_1 of Σ satisfying

$$\sin^2(\mathbf{w}_n, \mathbf{v}_1) \leq \epsilon \leq C \left(\frac{\mathcal{V} \log d}{(\lambda_1 - \lambda_2)^2} \cdot \frac{1}{n} + \left(\frac{2\beta}{n} \right)^{2 \log d} \right),$$

with probability greater than $3/4$. Here C is an absolute numerical constant.

The error above should be interpreted as being the sum of a $\mathcal{O}\left(\frac{1}{n}\right)$ higher order term and another $\mathcal{O}\left((2\beta/n)^{2\log d}\right)$ lower order term which is at most $o\left(\frac{1}{n^{\log d}}\right)$ (once $n > 4\beta^2$). In particular, this result shows that, up to an additive lower order term, one can match Theorem 1.1 with an asymptotic error of $\mathcal{O}\left(\frac{\mathcal{V}\log d}{(\lambda_1-\lambda_2)^2 n}\right)$ with constant probability. The lower order term has β which is the max of three parts: $\frac{\mathcal{M}\log d}{(\lambda_1-\lambda_2)}$, $\frac{\mathcal{V}\log^2 d}{(\lambda_1-\lambda_2)^2}$ and $\frac{\lambda_1^2 \log^2 d}{(\lambda_1-\lambda_2)^2}$. The first part, depending on \mathcal{M} , is exactly the same as what appears in Theorem 1.1. The second one, depending on \mathcal{V} has an additional $\log d$ factor over the first order term and is irrelevant once, say $n > 10\beta$. Notably, the third part, depending on λ_1^2 , does not appear in Theorem 1.1; it arises here entirely due to computational reasons: the setting allows only a *single linear-time pass* over the matrices, while Theorem 1.1 makes no such assumption. For instance, consider the case $\mathcal{V} = 0$ which means $\mathbf{A}_1 = \Sigma$. Matrix Bernstein tells us that one sample is sufficient to compute \mathbf{v}_1 . However, it is not evident how to compute it using a single pass over \mathbf{A}_1 . Note however, that the rate at which the lower order terms, i.e. $o\left(\frac{1}{n^{\log d}}\right)$, decrease is much better than $\mathcal{O}\left(1/n^2\right)$ guaranteed by Theorem 1.1.

In fact, this result also improves the asymptotic error rate obtained by Theorem 1.1. In particular, the following result shows that Oja's algorithm gets an asymptotic rate of $\mathcal{O}\left(\frac{\mathcal{V}}{(\lambda_1-\lambda_2)^2 n}\right)$ which is better than that of matrix Bernstein by a factor of $\mathcal{O}(\log d)$.²

Theorem 1.3. *Let the assumptions of Definition 1 hold. Suppose the step size sequence for Algorithm 1 is chosen to be $\eta_i = \frac{6}{(\lambda_1-\lambda_2)(\beta+i)}$, where*

$$\beta \triangleq 720 \max\left(\frac{\mathcal{M}}{(\lambda_1-\lambda_2)}, \frac{\mathcal{V} + \lambda_1^2}{(\lambda_1-\lambda_2)^2}\right).$$

Suppose $n > \beta^{1.2} d^{0.1}$. Then the output \mathbf{w}_n of Algorithm 1 is an ϵ -approximation to the top eigenvector \mathbf{v}_1 of Σ satisfying

$$\sin^2(\mathbf{w}_n, \mathbf{v}_1) \leq \epsilon \leq C \left(\frac{\mathcal{V}}{(\lambda_1-\lambda_2)^2} \cdot \frac{1}{n} + \frac{1}{n^2}\right),$$

with probability greater than $3/4$. Here C is an absolute numerical constant.

Note that Theorems 1.2 and 1.3 guarantee success probability of $3/4$. One way to boost the probability to $1 - \delta$, for some $\delta > 0$, is to run $\mathcal{O}(\log 1/\delta)$ copies of the algorithm, each with $3/4$ success probability and then output the geometric median of the solutions, which can be done in nearly linear time [17]. The details are omitted here.

Beyond the improved sample complexities we believe our analysis sheds light on the type of step sizes for which Oja's algorithm converges quickly and therefore illuminates how to efficiently perform streaming PCA. We note that we have essentially assumed an oracle which sets the step size sequence, and an important question is how to set the step size in a robust and data driven manner. Moreover, we believe that our analysis is fairly general and hope that it may be extended to make progress on analyzing the many variants of PCA that occur in both theory and in practice.

²A similar asymptotic result was recently obtained by [14]. However, their result requires an initial vector that is constant close to \mathbf{v}_1 , which itself is a difficult problem.

Algorithm	Error	$\mathcal{O}(d)$ space?
Oja's (this work, Theorem 4.1)	$\mathcal{O}\left(\frac{\nu}{(\lambda_1 - \lambda_2)^2} \cdot \frac{1}{n}\right)$	Yes
Matrix Bernstein + Wedin's theorem (Theorem 1.1)	$\mathcal{O}\left(\frac{\nu \log d}{(\lambda_1 - \lambda_2)^2} \cdot \frac{1}{n}\right)$	No
Alecton [9]	$\mathcal{O}\left(\frac{\nu d}{(\lambda_1 - \lambda_2)^2} \cdot \frac{\log n}{n}\right)$	Yes
Block Power Method [16]	$\mathcal{O}\left(\frac{\nu \lambda_1 \log d}{(\lambda_1 - \lambda_2)^3} \cdot \frac{\log n}{n}\right)$	Yes

Table 1: Asymptotic error guaranteed by various methods under assumptions of Definition 1 with at least constant probability, and ignoring constant factors. Recall that the error is defined as $\sin^2(\mathbf{w}, \mathbf{v}_1) = 1 - (\mathbf{w}^\top \mathbf{v}_1)^2$. Our analysis provides the optimal $1/n$ error decay rate as compared to Alecton and Block power method which obtain $\frac{\log n}{n}$. Moreover, our bound is $\mathcal{O}(d)$ tighter than that of Alecton [9] and $\mathcal{O}\left(\frac{\lambda_1}{\lambda_1 - \lambda_2}\right)$ tighter bound than that of Block Power Method [16]. The assumptions made in [9] for Alecton are different from our (more standard) assumption; we have optimized their bounds are optimized in our setting. See Section 1.1 for a concrete example where our analysis provides these improvements over [9, 16].

1.1 Comparison with Existing Results

Here we compare our sample complexity bounds with existing analyses of various methods. Recall that the error of the estimate \mathbf{w} is $\sin^2(\mathbf{w}, \mathbf{v}_1) = 1 - (\mathbf{w}^\top \mathbf{v}_1)^2$.

We consider three popular methods used for computing \mathbf{v}_1 . The first one is the batch method which computes largest eigenvector of empirical covariance and uses Wedin's theorem with matrix Bernstein inequality (cf. Theorem 1.1). The second method is Alecton, which is very similar to Oja's algorithm [9]. Finally, consider a block-power method (BPM) [16, 8] which divides samples into different blocks and applies power iteration to the empirical estimate from each block. See Table 1 for the comparison.

We stress that some of the results we compare to make different assumptions than Definition 1. The bounds stated for them are our best attempt to adapt their bounds in the setting of Definition 1 (which is quite standard). The next paragraph provides a simple example, which demonstrates the improvement in our result as compared to existing work.

Let $\mathbf{A}_i = \mathbf{x}_i \mathbf{x}_i^\top$, where $\mathbf{x}_i \in \mathbb{R}^d$ and $\mathbf{x}_i = \mathbf{e}_1$ with probability $1/d$ and $\mathbf{x}_i = \sigma \mathbf{e}_j$, $1 \leq j \leq d$ with probability $1/d$ where \mathbf{e}_j denotes the j^{th} standard basis vector and $\sigma < 1$. Note that $\Sigma = \mathbb{E}[\mathbf{A}_i] = \frac{(1-\sigma^2)}{d} \mathbf{e}_1 \mathbf{e}_1^\top + \frac{1}{d} \sigma^2 \mathbf{I}$, $\|\mathbf{A}_i\|_2 \leq 1$ for all i , and $\|\mathbb{E}[\mathbf{A}_i \mathbf{A}_i^\top]\|_2 \leq \frac{1}{d}$. Even for constant accuracy $\epsilon = \Omega(1)$, Theorem 1.2 tells us that $n = \mathcal{O}\left(\frac{d \log^2 d}{(1-\sigma^2)^2}\right)$ is sufficient. On the other hand, Theorem 1 of [9] requires $n = \mathcal{O}\left(\frac{d^2 \log^2 d}{(1-\sigma^2)^2}\right)$, while Theorem 2.4 of [16] requires $n = \mathcal{O}\left(\frac{d \log^2 d}{(1-\sigma^2)^3}\right)$. Asymptotically, as n becomes larger, our error scales as $\mathcal{O}\left(\frac{d}{(1-\sigma^2)^2} \cdot \frac{1}{n}\right)$ while that of [9] scales as $\mathcal{O}\left(\frac{d^2}{(1-\sigma^2)^2} \cdot \frac{\log n}{n}\right)$ and that of [16] scales as $\mathcal{O}\left(\frac{d}{(1-\sigma^2)^3} \cdot \frac{\log n}{n}\right)$. Combining matrix Bernstein and Wedin's theorems gives an asymptotic error of $\mathcal{O}\left(\frac{d \log d}{(1-\sigma^2)^2} \cdot \frac{1}{n}\right)$.

1.2 Additional Related Work

Existing results for computing largest eigenvector of a data covariance matrix using streaming samples can be divided into three broad settings: a) stochastic data, b) arbitrary sequence of data, c) regret bounds for arbitrary sequence of data.

Stochastic data: Here, the data is assumed to be sampled i.i.d. from a fixed distribution. The analysis of Oja’s algorithm as well as those of block power method and Alecon mentioned earlier are in this setting. [8] also obtained a result in the restricted spiked covariance model. [7] provides an analysis of a modification of Oja’s algorithm but with an extra $O(d^5)$ multiplicative factor compared to ours. [14] provides an algorithm based on shift and invert framework that obtains the same asymptotic error as ours. However, their algorithm requires warm start with a vector that is already constant close to the top eigenvector, which itself is a hard problem.

Arbitrary data: In this setting, each row of the data matrix is provided in an arbitrary order. Most of the existing methods here first compute a sketch of the matrix and use that to compute an estimate of the top eigenvector [18, 19, 20, 21, 22, 23]. However, a direct application of such techniques to the stochastic setting leads to sample complexity bounds which are larger by a multiplicative factor of $O(d)$ (ignoring other factors like variance etc). Finally, [24, 25, 14] also provide methods for eigenvector computation, but they require multiple passes over the data and hence do not apply to the streaming setting.

Regret bounds: Here, at each step the algorithm has to output an estimate \mathbf{w} of \mathbf{v}_1 for which we get reward of $\mathbf{w}^T \mathbf{A}_i \mathbf{w}$ and the goal is to minimize the regret w.r.t. \mathbf{v}_1 . The algorithms in this regime are mostly based on online convex optimization and applying them in our setting would again result in a loss of multiplicative $O(d)$. Moreover, typical algorithms in this setting are not memory efficient [26, 27].

1.3 Notation

Bold lowercase letters such as $\mathbf{u}, \mathbf{v}, \mathbf{w}$ are used to denote vectors and bold uppercase letters such as $\mathbf{A}, \mathbf{B}, \mathbf{C}$ to denote matrices. For symmetric matrices \mathbf{A} and \mathbf{B} , $\mathbf{A} \preceq \mathbf{B}$ denotes the condition that $\mathbf{x}^T \mathbf{A} \mathbf{x} \leq \mathbf{x}^T \mathbf{B} \mathbf{x}$ for all \mathbf{x} and define $\mathbf{B} \succeq \mathbf{A}$ analogously. A symmetric matrix \mathbf{A} is positive semidefinite if $\mathbf{A} \succeq \mathbf{0}$. For symmetric matrices \mathbf{A}, \mathbf{B} , define their inner product as $\langle \mathbf{A}, \mathbf{B} \rangle \triangleq \text{Tr}(\mathbf{A}^T \mathbf{B})$.

1.4 Paper Organization

The rest of this paper is organized as follows. Section 2 introduces basic mathematical facts used throughout the paper and also provides a proof of the error bound of the standard batch method (Theorem 1.1). Section 3 provides an overview of our approach to analyzing Oja’s algorithm and provides the main technical result of the paper. This technical result is used in Section 4 to prove the running time for Oja’s algorithm and to justify the choice of step size. Section 5 presents the proof of the main technical result. Section 6 concludes and mentions a few interesting future directions.

2 Preliminaries

The following basic inequalities regarding power series, the exponential, and PSD matrices are used throughout. The facts are summarized here:

Lemma 2.1 (Basic Inequalities). *The following are true:*

- $1 + x \leq \exp(x)$ for all x
- $1 + x \geq \exp(x - x^2)$ for all $x \geq 0$
- $\frac{1}{1+x} \leq \sum_{i=1}^{\infty} \frac{1}{(x+i)^2} \leq \frac{1}{x}$

- $\langle \mathbf{A}, \mathbf{B} \rangle \leq \langle \mathbf{A}, \mathbf{C} \rangle$ for PSD matrices $\mathbf{A}, \mathbf{B}, \mathbf{C}$ with $\mathbf{B} \preceq \mathbf{C}$
- $\text{Tr}(\mathbf{A}^\top \mathbf{B}) \leq \frac{1}{2} \text{Tr}(\mathbf{A}^\top \mathbf{A} + \mathbf{B}^\top \mathbf{B})$ for all matrices $\mathbf{A}, \mathbf{B} \in \mathbb{R}^{m \times n}$.

Proof. The first inequality follows from the Taylor expansion of $\exp(x)$. The second comes from $1 + 0 = \exp(0 - 0^2)$ and $\frac{d}{dx}(1 + x) \leq \frac{d}{dx} \exp(x - x^2)$ for $x \geq 0$. The third follows by considering upper and lower Riemann sums of $\int_{y=1}^{\infty} 1/(x + y)$. The fourth from the fact that since \mathbf{A} is PSD there is a matrix \mathbf{D} with $\mathbf{D}^\top \mathbf{D} = \mathbf{A}$ and therefore

$$\langle \mathbf{A}, \mathbf{B} \rangle = \text{Tr}(\mathbf{A}^\top \mathbf{B}) = \text{Tr}(\mathbf{D}\mathbf{B}\mathbf{D}^\top) \leq \text{Tr}(\mathbf{D}\mathbf{C}\mathbf{D}^\top) = \langle \mathbf{A}, \mathbf{C} \rangle .$$

The final follows from Cauchy Schwarz and Young's inequality, i.e. $x \cdot y \leq \frac{1}{2}(x^2 + y^2)$ as

$$\text{Tr}(\mathbf{B}^\top \mathbf{A}) = \sum_{i \in [n]} \mathbf{1}_i \mathbf{B}^\top \mathbf{A} \mathbf{1}_i \leq \sum_{i \in [n]} \|\mathbf{A} \mathbf{1}_i\|_2 \cdot \|\mathbf{B} \mathbf{1}_i\|_2 \leq \frac{1}{2} \sum_{i \in [n]} (\|\mathbf{A} \mathbf{1}_i\|_2^2 + \|\mathbf{B} \mathbf{1}_i\|_2^2)$$

□

The following is a matrix Bernstein based proof of the error bound of the batch method.

Proof of Theorem 1.1. Using Theorem 1.4 of [11], we have (w.p. $\geq 1 - \delta$):

$$\left\| \frac{1}{n} \sum_{i=1}^n \mathbf{A}_i - \Sigma \right\|_2 \leq 2 \cdot \max \left\{ \sqrt{\frac{\mathcal{V}}{n} \log \frac{d}{\delta}}, \frac{\mathcal{M}}{n} \log \frac{d}{\delta} \right\}. \quad (1)$$

Let $\hat{\mathbf{v}}$ be the top eigenvector of $\hat{\Sigma} = \frac{1}{n} \sum_{i=1}^n \mathbf{A}_i$. Using Wedin's theorem [13], implies:

$$\sin^2 \langle \mathbf{v}_1, \hat{\mathbf{v}} \rangle \leq \frac{\left\| \frac{1}{n} \sum_{i=1}^n \mathbf{A}_i - \Sigma \right\|_2^2}{|\lambda_1 - \lambda_2|^2}. \quad (2)$$

Theorem now follows by combining (1) and (2). □

3 Approach

Let us now describe the approach to analyze Oja's algorithm. We provide our main theorem regarding the convergence rate of Oja's algorithm and discuss how it is proved. The details of the proof are deferred to Section 5 and the use of the theorem to choose step sizes is in Section 4.

One of the primary difficulties in analyzing Oja's algorithm, or more broadly any algorithm for streaming PCA, is choosing a subtle potential function to analyze the method. If we try to analyze the progress of Oja's algorithm in every iteration i , by measuring the quality of \mathbf{w}_i , we run the risk that during the first few iterations of Oja's algorithm a step may actually yield a \mathbf{w}_{i+1} that is orthogonal to \mathbf{v}_i . If this happens, even in the typical best case, where all future samples are Σ itself, we would still fail to converge. In short, if we do not account for the randomness of \mathbf{w}_0 in our potential function then it is difficult to show that a rapidly convergent algorithm does not catastrophically fail.

Rather than analyzing the convergence of \mathbf{w}_i directly we instead analyze the convergence of Oja's algorithm as an operator on \mathbf{w}_0 . Oja's algorithm simply considers the matrix

$$\mathbf{B}_n \triangleq (\mathbf{I} + \eta_n \mathbf{A}_n)(\mathbf{I} + \eta_{n-1} \mathbf{A}_{n-1}) \cdots (\mathbf{I} + \eta_1 \mathbf{A}_1) \quad (3)$$

and outputs the normalized result of applying this matrix, \mathbf{B}_n , to the random initial vector, i.e.

$$\mathbf{w}_n = \frac{\mathbf{B}_n \mathbf{w}_0}{\|\mathbf{B}_n \mathbf{w}_0\|_2}. \quad (4)$$

Rather than analyze the improvement of \mathbf{w}_{n+1} over \mathbf{w}_n we analyze \mathbf{B}_{n+1} 's improvement over \mathbf{B}_n .

Another interpretation of (3) and (4) is that Oja's algorithm simply approximates \mathbf{v}_n by performing 1 step of the power method on the matrix \mathbf{B}_n . Fortunately, analyzing when 1 step of the power method succeeds is fairly straightforward as we show below:

Lemma 3.1 (One Step Power Method). *Let $\mathbf{B} \in \mathbb{R}^{d \times d}$, let $\tilde{\mathbf{v}} \in \mathbb{R}^d$ be a unit vector, and let $\tilde{\mathbf{V}}_\perp$ be a matrix whose columns form an orthonormal basis of the subspace orthogonal to $\tilde{\mathbf{v}}$. If $\mathbf{w} \in \mathbb{R}^d$ is chosen uniformly at random from the surface of the unit sphere then with probability at least $1 - \delta$*

$$\sin^2 \left(\tilde{\mathbf{v}}, \frac{\mathbf{B}\mathbf{w}}{\|\mathbf{B}\mathbf{w}\|_2} \right) = 1 - \left(\frac{\tilde{\mathbf{v}}^\top \mathbf{B}\mathbf{w}}{\|\mathbf{B}\mathbf{w}\|_2} \right)^2 \leq \frac{C \log(1/\delta)}{\delta} \frac{\text{Tr} \left(\tilde{\mathbf{V}}_\perp^\top \mathbf{B} \mathbf{B}^\top \tilde{\mathbf{V}}_\perp \right)}{\tilde{\mathbf{v}}^\top \mathbf{B} \mathbf{B}^\top \tilde{\mathbf{v}}}$$

where C is an absolute constant.

Proof. As \mathbf{w} is distributed uniformly over the sphere, we have: $\mathbf{w} = \mathbf{g}/\|\mathbf{g}\|_2$ where $\mathbf{g} \sim N(0, I)$. Consequently, with probability at least $1 - \delta$

$$\begin{aligned} 1 - \left(\frac{\tilde{\mathbf{v}}^\top \mathbf{B}\mathbf{w}}{\|\mathbf{B}\mathbf{w}\|_2} \right)^2 &= \frac{\mathbf{g}^\top \mathbf{B}^\top (\mathbf{I} - \tilde{\mathbf{v}}\tilde{\mathbf{v}}^\top) \mathbf{B}\mathbf{g}}{\mathbf{g}^\top \mathbf{B}^\top \mathbf{B}\mathbf{g}} \stackrel{\zeta_1}{\leq} \frac{C_1}{\delta} \frac{\mathbf{g}^\top \mathbf{B}^\top (\mathbf{I} - \tilde{\mathbf{v}}\tilde{\mathbf{v}}^\top) \mathbf{B}\mathbf{g}}{\tilde{\mathbf{v}}^\top \mathbf{B} \mathbf{B}^\top \tilde{\mathbf{v}}} \\ &\stackrel{\zeta_2}{\leq} \frac{C \log(1/\delta)}{\delta} \frac{\text{Tr} \left(\mathbf{B}^\top (\mathbf{I} - \tilde{\mathbf{v}}\tilde{\mathbf{v}}^\top) \mathbf{B} \right)}{\tilde{\mathbf{v}}^\top \mathbf{B} \mathbf{B}^\top \tilde{\mathbf{v}}}, \end{aligned}$$

where C_1 and C are absolute constants. ζ_1 follows as $\mathbf{g}^\top \mathbf{B}^\top \mathbf{B}\mathbf{g} \geq (\tilde{\mathbf{v}}^\top \mathbf{B}\mathbf{g})^2 \geq \frac{1}{C_1} \tilde{\mathbf{v}}^\top \mathbf{B} \mathbf{B}^\top \tilde{\mathbf{v}}$ where the second inequality follows from the fact that $\tilde{\mathbf{v}}^\top \mathbf{B}\mathbf{g}$ is a Gaussian random variable with variance $\|\mathbf{B}^\top \tilde{\mathbf{v}}\|_2^2$. Similarly, ζ_2 follows from the fact that $\mathbf{g}^\top \mathbf{B}^\top (\mathbf{I} - \tilde{\mathbf{v}}\tilde{\mathbf{v}}^\top) \mathbf{B}\mathbf{g}$ is a χ^2 random variable with $\text{Tr} \left(\mathbf{B}^\top (\mathbf{I} - \tilde{\mathbf{v}}\tilde{\mathbf{v}}^\top) \mathbf{B} \right)$ degrees of freedom. \square

This lemma makes our goal clear. To show that Oja's algorithm succeeds we simply need to show that with constant probability $\mathbf{v}_1^\top \mathbf{B}_n \mathbf{B}_n^\top \mathbf{v}_1$ is relatively large and $\text{Tr} \left(\mathbf{V}_\perp \mathbf{B}_n \mathbf{B}_n^\top \mathbf{V}_\perp \right)$ is relatively small, where \mathbf{V}_\perp is a matrix whose columns form an orthonormal basis of the subspace orthogonal to \mathbf{v}_1 . This immediately alleviates the issues of catastrophic failure that plagued analyzing \mathbf{w}_n . So long as we pick η_i sufficiently small, i.e. $\eta_i = O(1/\max\{\mathcal{M}, \lambda_1\})$ then $\mathbf{I} + \eta_i \mathbf{A}_i$ is invertible. In this case $\mathbf{B}_n \mathbf{B}_n^\top$ is invertible and $\mathbf{v}_1^\top \mathbf{B}_n \mathbf{B}_n^\top \mathbf{v}_1 > 0$. In short, so long as we pick η_i sufficiently small the quantity we wish to bound $\text{Tr} \left(\mathbf{V}_\perp \mathbf{B}_n \mathbf{B}_n^\top \mathbf{V}_\perp \right) / \mathbf{v}_1^\top \mathbf{B}_n \mathbf{B}_n^\top \mathbf{v}_1$ is always finite.

To actually bound $\mathbf{v}_1^\top \mathbf{B}_n \mathbf{B}_n^\top \mathbf{v}_1$ and $\text{Tr} \left(\mathbf{V}_\perp \mathbf{B}_n \mathbf{B}_n^\top \mathbf{V}_\perp \right)$ we split the analysis into several parts in Section 5. First, we show that $\mathbb{E} \left[\text{Tr} \left(\mathbf{V}_\perp \mathbf{B}_n \mathbf{B}_n^\top \mathbf{V}_\perp \right) \right]$ is small, which implies by Markov's inequality that $\text{Tr} \left(\mathbf{V}_\perp \mathbf{B}_n \mathbf{B}_n^\top \mathbf{V}_\perp \right)$ is small with constant probability. Then, we show that $\mathbb{E} \mathbf{v}_1^\top \mathbf{B}_n \mathbf{B}_n^\top \mathbf{v}_1$ is large and that $\text{Var} \left[\mathbf{v}_1^\top \mathbf{B}_n \mathbf{B}_n^\top \mathbf{v}_1 \right]$ is small. By Chebyshev's inequality this implies that $\mathbf{v}_1^\top \mathbf{B}_n \mathbf{B}_n^\top \mathbf{v}_1$ is large with constant probability. Putting these together we achieve the main technical result regarding the analysis of Oja's method. Once we devise this roadmap, the proof is fairly straightforward.

Theorem 3.1 (Oja’s Algorithm Convergence Rate). *Let $\delta > 0$ and step sizes $\eta_i \leq \frac{1}{4 \cdot \max\{M, \lambda_1\}}$. The output \mathbf{w}_n of Algorithm 1 is an ϵ -approximation to \mathbf{v}_1 with probability at least $1 - \delta$ where*

$$\epsilon \leq \frac{1}{Q} \exp \left(5\bar{\mathcal{V}} \sum_{i \in [n]} \eta_i^2 \right) \left(d \cdot \exp \left(-2(\lambda_1 - \lambda_2) \sum_{i \in [n]} \eta_i \right) + \mathcal{V} \sum_{i=1}^n \eta_i^2 \exp \left(- \sum_{j=i+1}^n 2\eta_j (\lambda_1 - \lambda_2) \right) \right)$$

where $Q \triangleq \frac{\delta^2}{C \log(1/\delta)} \left(1 - \frac{1}{\sqrt{\delta}} \sqrt{\exp(18\bar{\mathcal{V}} \sum_{i=1}^n \eta_i^2) - 1} \right)$, $\bar{\mathcal{V}} \triangleq \mathcal{V} + \lambda_1^2$, and C is an absolute constant.

Theorem 3.1 is proved in Section 5. Theorem 3.1 serves as the basis for our results regarding Oja’s algorithm. In the next section we show how to use this theorem to choose step sizes and achieve the main results of this paper.

4 Main Results

Theorem 3.1, from the previous section, leads to our main results, provided here. The theorem and proof are below and essentially consist of choosing appropriate parameters to efficiently apply Theorem 3.1. Once we have this theorem, Theorems 1.2 and 1.3 follow by choosing $\alpha = \log d$ and $\alpha = 6$ respectively.

Theorem 4.1. *Fix any $\delta > 0$ and suppose the step sizes are set to $\eta_t = \frac{\alpha}{(\lambda_1 - \lambda_2)(\beta + t)}$ for $\alpha > \frac{1}{2}$ and*

$$\beta \triangleq 20 \max \left(\frac{\mathcal{M}\alpha}{(\lambda_1 - \lambda_2)}, \frac{(\mathcal{V} + (\lambda_1)^2) \alpha^2}{(\lambda_1 - \lambda_2)^2 \log \left(1 + \frac{\delta}{100} \right)} \right).$$

Suppose the number of samples $n > \beta$. Then the output \mathbf{w}_n of Algorithm 1 satisfies:

$$1 - (\mathbf{w}_n^\top \mathbf{v}_1)^2 \leq \frac{C \log(1/\delta)}{\delta^2} \left(d \left(\frac{\beta}{n} \right)^{2\alpha} + \frac{\alpha^2 \mathcal{V}}{(2\alpha - 1)(\lambda_1 - \lambda_2)^2} \cdot \frac{1}{n} \right),$$

with probability at least $1 - \delta$. Here C is an absolute numerical constant.

Proof. Recall that Theorem 3.1 gives a bound of

$$\frac{1}{Q} \exp \left(5\bar{\mathcal{V}} \sum_{i \in [n]} \eta_i^2 \right) \left(d \cdot \exp \left(-2(\lambda_1 - \lambda_2) \sum_{i \in [n]} \eta_i \right) + \mathcal{V} \sum_{i=1}^n \eta_i^2 \exp \left(- \sum_{j=i+1}^n 2\eta_j (\lambda_1 - \lambda_2) \right) \right) \quad (5)$$

where $Q \triangleq \frac{\delta^2}{C \log(1/\delta)} \left(1 - \frac{1}{\sqrt{\delta}} \sqrt{\exp(18\bar{\mathcal{V}} \sum_{i=1}^n \eta_i^2) - 1} \right)$. Since $\eta_i = \frac{\alpha}{(\lambda_1 - \lambda_2)(\beta + i)}$, we have $\sum_{i \in [n]} \eta_i^2 \leq \frac{\alpha^2}{(\lambda_1 - \lambda_2)^2 \beta}$ and by our assumption that $\frac{\bar{\mathcal{V}} \alpha^2}{(\lambda_1 - \lambda_2)^2 \beta} \leq \frac{1}{18} \log \left(1 + \frac{\delta}{100} \right)$, we have:

$$\exp \left(18\bar{\mathcal{V}} \sum_{i \in [n]} \eta_i^2 \right) \leq \sqrt{2} \quad \Rightarrow \quad Q \geq \frac{\delta^2}{C \log(1/\delta)}. \quad (6)$$

Moreover, since $\sum_{i \in [n]} \eta_i \geq \frac{\alpha}{\lambda_1 - \lambda_2} \log(1 + n/\beta)$, we have

$$\exp\left(-2(\lambda_1 - \lambda_2) \sum_{i \in [n]} \eta_i\right) \leq \left(\frac{\beta}{\beta + n}\right)^{2\alpha}. \quad (7)$$

Note that $\sum_{j=i+1}^n \eta_j \leq \frac{\alpha}{\lambda_1 - \lambda_2} \log \frac{n+\beta+1}{i+\beta+1}$. Moreover, as $\alpha > 1/2$, we have:

$$\begin{aligned} & \sum_{i=1}^n \eta_i^2 \exp\left(-2(\lambda_1 - \lambda_2) \sum_{j=i+1}^n \eta_j\right) \\ & \leq \frac{\alpha^2}{(\lambda_1 - \lambda_2)^2} \sum_{i=1}^n \frac{1}{(\beta + i)^2} \exp\left(2\alpha \log \frac{i + \beta + 1}{n + \beta + 1}\right), \\ & \leq \frac{(\beta + 1)^2}{\beta^2} \cdot \frac{\alpha^2}{(\lambda_1 - \lambda_2)^2 (n + \beta + 1)^{2\alpha}} \cdot \sum_{i=1}^n (i + \beta + 1)^{2\alpha - 2}, \\ & \leq \frac{2\alpha^2}{(2\alpha - 1)(\lambda_1 - \lambda_2)^2 (n + \beta + 1)} \quad \left(\text{since } \alpha > 1/2 \text{ and } \sum_{i=1}^n i^\gamma \leq n^{\gamma+1}/(\gamma + 1) \forall \gamma > -1\right). \end{aligned} \quad (8)$$

Substituting (6), (7) and (8) into (5) proves the theorem. \square

5 Bounding the Convergence of Oja's Algorithm

In this section, we present a detailed proof of Theorem 3.1. The proof follows the approach outlined in Section 3 and uses the notation of that section, i.e.

- We let $\mathbf{B}_n \triangleq (\mathbf{I} + \eta_n \mathbf{A}_n) \cdots (\mathbf{I} + \eta_1 \mathbf{A}_1)$ with $\mathbf{B}_0 \triangleq \mathbf{I}$
- We let $\bar{\mathcal{V}} \triangleq \mathcal{V} + \lambda_1^2$
- We let $\mathbf{V}_\perp \in \mathbb{R}^{d \times d-1}$ denote a matrix whose columns form an orthonormal basis for the subspace orthogonal to \mathbf{v}_1 .

We first provide several technical lemmas bounding the expected behavior of \mathbf{B}_n and ultimately use these lemmas to prove Theorem 3.1. We begin with a straightforward lemma bounding the rate of increase of $\mathbb{E}[\mathbf{B}_t \mathbf{B}_t^\top]$ in spectral norm.

Lemma 5.1. *For all $t \geq 0$ and $\eta_i \geq 0$ we have*

$$\left\| \mathbb{E}[\mathbf{B}_t \mathbf{B}_t^\top] \right\|_2 \leq \exp\left(\sum_{i \in [t]} 2\eta_i \lambda_1 + \eta_i^2 \bar{\mathcal{V}}\right).$$

Proof. Let $\alpha_t \triangleq \|\mathbb{E}[\mathbf{B}_t \mathbf{B}_t^\top]\|_2$, i.e., $\mathbb{E}[\mathbf{B}_t \mathbf{B}_t^\top] \preceq \alpha_t \mathbf{I}$. For all $t > 0$,

$$\begin{aligned} \mathbb{E}[\mathbf{B}_t \mathbf{B}_t^\top] &= \mathbb{E}\left[(\mathbf{I} + \eta_t \mathbf{A}_t) \mathbf{B}_{t-1} \mathbf{B}_{t-1}^\top (\mathbf{I} + \eta_t \mathbf{A}_t)^\top\right] \preceq \alpha_{t-1} \mathbb{E}\left[(\mathbf{I} + \eta_t \mathbf{A}_t)(\mathbf{I} + \eta_t \mathbf{A}_t^\top)\right], \\ &= \alpha_{t-1} \mathbb{E}\left[\mathbf{I} + \eta_t \mathbf{A}_t + \eta_t \mathbf{A}_t^\top + \eta_t^2 \mathbf{A}_t \mathbf{A}_t^\top\right] \preceq \alpha_{t-1} \left[\mathbf{I} + 2\eta_t \boldsymbol{\Sigma} + \eta_t^2 (\boldsymbol{\Sigma}^2 + \mathbf{V}\mathbf{I})\right], \end{aligned} \quad (9)$$

where the last inequality follows from $\mathbb{E}[A_t] = \Sigma$ and,

$$\mathbb{E}[\mathbf{A}_t \mathbf{A}_t^\top] = \Sigma^2 + \mathbb{E}[(\mathbf{A}_t - \Sigma)(\mathbf{A}_t - \Sigma)^\top] \preceq \Sigma^2 + V\mathbf{I}.$$

Using (9) along with $\|\mathbb{E}[\mathbf{B}_t \mathbf{B}_t^\top]\|_2 = \alpha_t$, $\Sigma \preceq \lambda_1 \mathbf{I}$, and $\Sigma^2 \preceq \lambda_1^2 \mathbf{I}$, we have for $\forall t > 0$:

$$\alpha_t \leq (1 + 2\eta_t \lambda_1 + \eta_t^2 (\lambda_1^2 + V)) \alpha_{t-1}.$$

The result follows by using induction along with $\alpha_0 = 1$ and $1 + x \leq e^x$. \square

Using Lemma 5.1 we next bound the expected value of $\text{Tr}(\mathbf{V}_\perp^\top \mathbf{B}_n \mathbf{B}_n^\top \mathbf{V}_\perp)$. Ultimately this will allow us to bound the value $\text{Tr}(\mathbf{V}_\perp^\top \mathbf{B}_n \mathbf{B}_n^\top \mathbf{V}_\perp)$ with by Markov's inequality.

Lemma 5.2. *For all $t \geq 0$ and $\eta_i \leq \frac{1}{\lambda_1}$ the following holds*

$$\mathbb{E}[\text{Tr}(\mathbf{V}_\perp^\top \mathbf{B}_t \mathbf{B}_t^\top \mathbf{V}_\perp)] \leq \exp\left(\sum_{j \in [t]} 2\eta_j \lambda_2 + \eta_j^2 \bar{\mathcal{V}}\right) \cdot \left(d + \mathcal{V} \sum_{i=1}^t \eta_i^2 \exp\left(\sum_{j \in [i]} 2\eta_j (\lambda_1 - \lambda_2)\right)\right).$$

Proof. Let $\alpha_t \triangleq \mathbb{E}[\text{Tr}(\mathbf{V}_\perp^\top \mathbf{B}_t \mathbf{B}_t^\top \mathbf{V}_\perp)]$. We first simplify α_t as follows:

$$\alpha_t = \left\langle \mathbb{E}[\mathbf{B}_t \mathbf{B}_t^\top], \mathbf{V}_\perp \mathbf{V}_\perp^\top \right\rangle = \left\langle \mathbb{E}[\mathbf{B}_{t-1} \mathbf{B}_{t-1}^\top], \mathbb{E}[(\mathbf{I} + \eta_t \mathbf{A}_t) \mathbf{V}_\perp \mathbf{V}_\perp^\top (\mathbf{I} + \eta_t \mathbf{A}_t^\top)] \right\rangle. \quad (10)$$

Recall that $\mathbb{E}[\mathbf{A}_t] = \Sigma$. Now, the second term on the right hand side can be bounded as follows:

$$\begin{aligned} & \mathbb{E}[(\mathbf{I} + \eta_t \mathbf{A}_t) \mathbf{V}_\perp \mathbf{V}_\perp^\top (\mathbf{I} + \eta_t \mathbf{A}_t^\top)], \\ &= \mathbf{V}_\perp \mathbf{V}_\perp^\top + \eta_t \Sigma \mathbf{V}_\perp \mathbf{V}_\perp^\top + \eta_t \mathbf{V}_\perp \mathbf{V}_\perp^\top \Sigma + \eta_t^2 \mathbb{E}[\mathbf{A}_t \mathbf{V}_\perp \mathbf{V}_\perp^\top \mathbf{A}_t^\top], \\ &= \mathbf{V}_\perp \mathbf{V}_\perp^\top + \eta_t \Sigma \mathbf{V}_\perp \mathbf{V}_\perp^\top + \eta_t \mathbf{V}_\perp \mathbf{V}_\perp^\top \Sigma + \eta_t^2 \Sigma \mathbf{V}_\perp \mathbf{V}_\perp^\top \Sigma + \eta_t^2 \mathbb{E}[(\mathbf{A}_t - \Sigma) \mathbf{V}_\perp \mathbf{V}_\perp^\top (\mathbf{A}_t - \Sigma)^\top], \\ &\stackrel{\zeta_1}{\leq} \mathbf{V}_\perp \mathbf{V}_\perp^\top + 2\eta_t \lambda_2 \mathbf{V}_\perp \mathbf{V}_\perp^\top + \eta_t^2 \lambda_2^2 \mathbf{V}_\perp \mathbf{V}_\perp^\top + \eta_t^2 \mathbb{E}[(\mathbf{A}_t - \Sigma) (\mathbf{A}_t - \Sigma)^\top], \\ &\stackrel{\zeta_2}{\leq} (1 + 2\eta_t \lambda_2 + \eta_t^2 \lambda_2^2) \mathbf{V}_\perp \mathbf{V}_\perp^\top + \eta_t^2 \mathcal{V} \mathbf{I} = (1 + 2\eta_t \lambda_2 + \eta_t^2 \lambda_2^2 + \eta_t^2 \mathcal{V}) \mathbf{V}_\perp \mathbf{V}_\perp^\top + \eta_t^2 \mathcal{V} \cdot \mathbf{v}_1 \mathbf{v}_1^\top, \end{aligned}$$

where ζ_1 follows from the fact that \mathbf{V}_\perp is orthogonal to \mathbf{v}_1 and ζ_2 follows from definition of \mathcal{V} .

Plugging the above into (10), we get for all $t \geq 1$,

$$\begin{aligned} \alpha_t &\leq (1 + 2\eta_t \lambda_2 + \eta_t^2 (\lambda_2^2 + \mathcal{V})) \left\langle \mathbb{E}[\mathbf{B}_{t-1} \mathbf{B}_{t-1}^\top], \mathbf{V}_\perp \mathbf{V}_\perp^\top \right\rangle + \eta_t^2 \mathcal{V} \left\langle \mathbb{E}[\mathbf{B}_{t-1} \mathbf{B}_{t-1}^\top], \mathbf{v}_1 \mathbf{v}_1^\top \right\rangle, \\ &\leq (1 + 2\eta_t \lambda_2 + \eta_t^2 \bar{\mathcal{V}}) \alpha_{t-1} + \eta_t^2 \mathcal{V} \left\| \mathbb{E}[\mathbf{B}_{t-1} \mathbf{B}_{t-1}^\top] \right\|_2, \\ &\leq \exp(2\eta_t \lambda_2 + \eta_t^2 \bar{\mathcal{V}}) \alpha_{t-1} + \eta_t^2 \mathcal{V} \exp\left(\sum_{i \in [t-1]} \eta_i \lambda_1 + \eta_i^2 \bar{\mathcal{V}}\right), \end{aligned}$$

where the last inequality follows from $1 + x \leq e^x$ and using Lemma 5.1.

Recurring the above inequality, we obtain

$$\begin{aligned}\alpha_t &\leq \sum_{i \in [t]} \eta_i^2 \mathcal{V} \exp \left(\sum_{j=i+1}^t 2\eta_j \lambda_2 + \eta_j^2 \bar{\mathcal{V}} \right) \exp \left(\sum_{j \in [i]} 2\eta_j \lambda_1 + \eta_j^2 \bar{\mathcal{V}} \right) + \exp \left(\sum_{j \in [t]} 2\eta_j \lambda_2 + \eta_j^2 \bar{\mathcal{V}} \right) \alpha_0, \\ &\leq \exp \left(\sum_{j \in [t]} 2\eta_j \lambda_2 + \eta_j^2 \bar{\mathcal{V}} \right) \left(\alpha_0 + \mathcal{V} \sum_{i=1}^t \eta_i^2 \exp \left(\sum_{j \in [i]} 2\eta_j (\lambda_1 - \lambda_2) + \eta_j^2 \bar{\mathcal{V}} \right) \right)\end{aligned}$$

Since $\mathbf{B}_0 = \mathbf{I}$ we see that $\alpha_0 = d - 1 \leq d$. Using that $\eta_i \leq \frac{1}{\lambda_1} \leq \frac{1}{\lambda_2}$ completes the proof. \square

Next we provide the lemmas that will allow us to lower bound $\mathbf{v}_1^\top \mathbf{B}_t \mathbf{B}_t^\top \mathbf{v}_1$. In Lemma 5.3 we lower bound $\mathbb{E} [\mathbf{v}_1^\top \mathbf{B}_t \mathbf{B}_t^\top \mathbf{v}_1]$ and in Lemma 5.4 we upper bound $\text{Var} [\mathbf{v}_1^\top \mathbf{B}_t \mathbf{B}_t^\top \mathbf{v}_1]$. Ultimately, the lower bound follows using Chebyshev's inequality.

Lemma 5.3. *For all $t \geq 0$ and $\eta_i \geq 0$ we have*

$$\mathbb{E} [\mathbf{v}_1^\top \mathbf{B}_t \mathbf{B}_t^\top \mathbf{v}_1] \geq \exp \left(\sum_{i \in [t]} 2\eta_i \lambda_1 - 4\eta_i^2 \lambda_1^2 \right)$$

If we further assume that $\eta_i \leq \frac{1}{4 \cdot \max\{\lambda_1, M\}}$ then $\mathbb{E} [\mathbf{v}_1^\top \mathbf{B}_t \mathbf{B}_t^\top \mathbf{v}_1] \geq \exp(\lambda_1 \sum_{i \in [t]} \eta_i)$.

Proof. Let $\beta_t \triangleq \mathbb{E} [\mathbf{v}_1^\top \mathbf{B}_t \mathbf{B}_t^\top \mathbf{v}_1]$. Since $\mathbf{B}_t = (\mathbf{I} + \eta_t \mathbf{A}_t) \mathbf{B}_{t-1}$, we can bound β_t as

$$\begin{aligned}\beta_t &= \left\langle \mathbb{E} [\mathbf{B}_{t-1} \mathbf{B}_{t-1}^\top], \mathbb{E} [(\mathbf{I} + \eta_t \mathbf{A}_t) \mathbf{v}_1 \mathbf{v}_1^\top (\mathbf{I} + \eta_t \mathbf{A}_t^\top)] \right\rangle \\ &= \left\langle \mathbb{E} [\mathbf{B}_{t-1} \mathbf{B}_{t-1}^\top], \mathbf{v}_1 \mathbf{v}_1^\top + \eta_t \boldsymbol{\Sigma} \mathbf{v}_1 \mathbf{v}_1^\top + \eta_t \mathbf{v}_1 \mathbf{v}_1^\top \boldsymbol{\Sigma} + \eta_t^2 \mathbb{E} [\mathbf{A}_t \mathbf{v}_1 \mathbf{v}_1^\top \mathbf{u}^{*\top} \mathbf{A}_t^\top] \right\rangle \\ &\geq \left\langle \mathbb{E} [\mathbf{B}_{t-1} \mathbf{B}_{t-1}^\top], \mathbf{v}_1 \mathbf{v}_1^\top + \lambda_1 \eta_t \mathbf{v}_1 \mathbf{v}_1^\top + \lambda_1 \eta_t \mathbf{v}_1 \mathbf{v}_1^\top \right\rangle.\end{aligned}$$

Consequently $\beta_t \geq (1 + 2\eta_t \lambda_1) \beta_{t-1}$. Furthermore, $\mathbf{B}_0 = \mathbf{I}$ and hence $\beta_0 = \|\mathbf{v}_1\|_2^2 = 1$. Proceeding by induction and using that $1 + x \geq \exp(x - x^2)$ for all $x \geq 0$ finishes the proof. \square

Lemma 5.4. *For $t \geq 0$ suppose that $\eta_i \leq \frac{1}{4 \cdot \max\{\lambda_1, M\}}$ for all $i \in [t]$ then.*

$$\mathbb{E} \left[\left(\mathbf{v}_1^\top \mathbf{B}_t \mathbf{B}_t^\top \mathbf{v}_1 \right)^2 \right] \leq \exp \left(\sum_{i \in [t]} 4\eta_i \lambda_1 + 10\eta_i^2 \bar{\mathcal{V}} \right)$$

Proof. Let $\mathbf{W}_{t,s} \triangleq (\mathbf{I} + \eta_t \mathbf{A}_t) \cdots (\mathbf{I} + \eta_{s+1} \mathbf{A}_{s+1})$ and $\gamma_s \triangleq \mathbb{E} \left[\left(\mathbf{v}_1^\top \mathbf{W}_{t,s} \mathbf{W}_{t,s}^\top \mathbf{v}_1 \right)^2 \right]$. Note that $\mathbf{W}_{t,t} = \mathbf{B}_t$ and $\gamma_t = \mathbb{E} [\mathbf{v}_1^\top \mathbf{B}_t \mathbf{B}_t^\top \mathbf{v}_1]$. Now,

$$\begin{aligned}\gamma_t &= \text{Tr} \left(\mathbb{E} \left[\mathbf{W}_{t,t}^\top \mathbf{v}_1 \mathbf{v}_1^\top \mathbf{W}_{t,t} \mathbf{W}_{t,t}^\top \mathbf{v}_1 \mathbf{v}_1^\top \mathbf{W}_{t,t} \right] \right) \\ &= \text{Tr} \left(\mathbb{E} \left[(\mathbf{I} + \eta_1 \mathbf{A}_1^\top) \mathbf{W}_{t,t-1}^\top \mathbf{v}_1 \mathbf{v}_1^\top \mathbf{W}_{t,t-1} (\mathbf{I} + \eta_1 \mathbf{A}_1) (\mathbf{I} + \eta_1 \mathbf{A}_1^\top) \mathbf{W}_{t,t-1}^\top \mathbf{v}_1 \mathbf{v}_1^\top \mathbf{W}_{t,t-1} (\mathbf{I} + \eta_1 \mathbf{A}_1) \right] \right) \\ &= \text{Tr} \left(\mathbb{E} \left[(\mathbf{I} + \eta_1 \mathbf{A}_1^\top) \mathbf{G}_{t-1} (\mathbf{I} + \eta_1 \mathbf{A}_1) (\mathbf{I} + \eta_1 \mathbf{A}_1^\top) \mathbf{G}_{t-1} (\mathbf{I} + \eta_1 \mathbf{A}_1) \right] \right),\end{aligned}\tag{11}$$

where $\mathbf{G}_{t-1} \triangleq \mathbf{W}_{t,t-1}^\top \mathbf{v}_1 \mathbf{v}_1^\top \mathbf{W}_{t,t-1}$. In order to bound the above quantity, we first bound the above expression for an arbitrary $\mathbf{G}_{t-1} \equiv \mathbf{G}$. We then take an expectation over only \mathbf{A}_1 and then finally take an expectation over \mathbf{G}_{t-1} . That is, for an arbitrary fixed symmetric matrix \mathbf{G} , we have:

$$\begin{aligned}
& \text{Tr} \left(\mathbb{E} \left[\left(\mathbf{I} + \eta_1 \mathbf{A}_1^\top \right) \mathbf{G} \left(\mathbf{I} + \eta_1 \mathbf{A}_1 \right) \left(\mathbf{I} + \eta_1 \mathbf{A}_1^\top \right) \mathbf{G} \left(\mathbf{I} + \eta_1 \mathbf{A}_1 \right) \right] \right) \\
&= \text{Tr} \left(\mathbb{E} \left[\left(\mathbf{G} + \eta_1 \mathbf{A}_1^\top \mathbf{G} + \eta_1 \mathbf{G} \mathbf{A}_1 + \eta_1^2 \mathbf{A}_1^\top \mathbf{G} \mathbf{A}_1 \right)^2 \right] \right) \\
&= \text{Tr} \left(\mathbf{G}^2 + \eta_1 \mathbb{E} \left[\mathbf{A}_1^\top \right] \mathbf{G}^2 + \eta_1 \mathbf{G}^2 \mathbb{E} \left[\mathbf{A}_1 \right] + \eta_1 \mathbf{G} \left(\mathbb{E} \left[\mathbf{A}_1 \right] + \mathbb{E} \left[\mathbf{A}_1^\top \right] \right) \mathbf{G} \right. \\
&\quad + \eta_1^2 \mathbb{E} \left[\mathbf{A}_1^\top \mathbf{G} \mathbf{A}_1 \mathbf{G} \right] + \eta_1^2 \mathbb{E} \left[\mathbf{A}_1^\top \mathbf{G} \mathbf{A}_1^\top \mathbf{G} \right] + \eta_1^2 \mathbb{E} \left[\mathbf{G} \mathbf{A}_1 \mathbf{G} \mathbf{A}_1 \right] + \eta_1^2 \mathbb{E} \left[\mathbf{G} \mathbf{A}_1^\top \mathbf{G} \mathbf{A}_1 \right] \\
&\quad + \eta_1^2 \mathbf{G} \mathbb{E} \left[\mathbf{A}_1 \mathbf{A}_1^\top \right] \mathbf{G} + \eta_1^2 \mathbb{E} \left[\mathbf{A}_1^\top \mathbf{G}^2 \mathbf{A}_1 \right] + \eta_1^3 \mathbb{E} \left[\mathbf{A}_1^\top \mathbf{G} \left(\mathbf{A}_1 + \mathbf{A}_1^\top \right) \mathbf{G} \mathbf{A}_1 \right] \\
&\quad \left. + \eta_1^3 \mathbb{E} \left[\mathbf{A}_1^\top \mathbf{G} \mathbf{A}_1 \mathbf{A}_1^\top \mathbf{G} \right] + \eta_1^3 \mathbb{E} \left[\mathbf{G} \mathbf{A}_1 \mathbf{A}_1^\top \mathbf{G} \mathbf{A}_1 \right] + \eta_1^4 \mathbb{E} \left[\mathbf{A}_1^\top \mathbf{G} \mathbf{A}_1 \mathbf{A}_1^\top \mathbf{G} \mathbf{A}_1 \right] \right) \\
&= \text{Tr} \left(\mathbf{G}^2 \right) + 4\eta_1 \text{Tr} \left(\Sigma \mathbf{G}^2 \right) + 2\eta_1^2 \text{Tr} \left(\mathbb{E} \left[\mathbf{A}_1 \mathbf{A}_1^\top \right] \mathbf{G}^2 \right) + \eta_1^2 \text{Tr} \left(\mathbb{E} \left[\mathbf{A}_1^\top \mathbf{G} \mathbf{A}_1 \mathbf{G} \right] \right) \\
&\quad + \eta_1^2 \text{Tr} \left(\mathbb{E} \left[\mathbf{A}_1^\top \mathbf{G} \mathbf{A}_1^\top \mathbf{G} \right] \right) + \eta_1^2 \text{Tr} \left(\mathbb{E} \left[\mathbf{G} \mathbf{A}_1 \mathbf{G} \mathbf{A}_1 \right] \right) + \eta_1^2 \text{Tr} \left(\mathbb{E} \left[\mathbf{G} \mathbf{A}_1^\top \mathbf{G} \mathbf{A}_1 \right] \right) \\
&\quad + 2\eta_1^3 \text{Tr} \left(\mathbb{E} \left[\mathbf{A}_1^\top \mathbf{G} \left(\mathbf{A}_1 + \mathbf{A}_1^\top \right) \mathbf{G} \mathbf{A}_1 \right] \right) + \eta_1^4 \text{Tr} \left(\mathbb{E} \left[\mathbf{A}_1^\top \mathbf{G} \mathbf{A}_1 \mathbf{A}_1^\top \mathbf{G} \mathbf{A}_1 \right] \right) \tag{12}
\end{aligned}$$

We now bound the various terms above as follows. Each of the second order terms can be bounded using Lemma 2.1 as follows:

$$\begin{aligned}
\mathbb{E} \left[\text{Tr} \left(\mathbf{A}_1^\top \mathbf{G} \mathbf{A}_1 \mathbf{G} \right) \right] &\leq \frac{1}{2} \mathbb{E} \left[\left\| \mathbf{A}_1^\top \mathbf{G} \right\|_F^2 + \left\| \mathbf{A}_1 \mathbf{G} \right\|_F^2 \right] \\
&= \frac{1}{2} \left(\text{Tr} \left(\mathbf{G} \mathbb{E} \left[\mathbf{A}_1 \mathbf{A}_1^\top \right] \mathbf{G} + \mathbf{G} \mathbb{E} \left[\mathbf{A}_1^\top \mathbf{A}_1 \right] \mathbf{G} \right) \right) \leq (\nu + \lambda_1^2) \text{Tr} \left(\mathbf{G}^2 \right). \tag{13}
\end{aligned}$$

The third order terms can be bounded as follows:

$$\begin{aligned}
\mathbb{E} \left[\text{Tr} \left(\mathbf{A}_1^\top \mathbf{G} \mathbf{A}_1 \mathbf{G} \mathbf{A}_1 \right) \right] &\leq \mathbb{E} \left[\left\| \mathbf{A}_1 \right\|_2 \text{Tr} \left(\mathbf{A}_1^\top \mathbf{G} \mathbf{G} \mathbf{A}_1 \right) \right] \\
&\leq (\mathcal{M} + \lambda_1) \text{Tr} \left(\mathbf{G} \mathbb{E} \left[\mathbf{A}_1 \mathbf{A}_1^\top \right] \mathbf{G} \right) \leq (\mathcal{M} + \lambda_1) \bar{\nu} \cdot \text{Tr} \left(\mathbf{G}^2 \right). \tag{14}
\end{aligned}$$

where we used the assumption that $\left\| \mathbf{A}_1 \right\|_2 \leq \left\| \mathbf{A}_1 - \Sigma \right\|_2 + \left\| \Sigma \right\|_2 \leq \mathcal{M} + \lambda_1$ with probability 1. Finally the fourth order term can be bounded as

$$\text{Tr} \left(\mathbb{E} \left[\mathbf{A}_1^\top \mathbf{G} \mathbf{A}_1 \mathbf{A}_1^\top \mathbf{G} \mathbf{A}_1 \right] \right) \leq (\mathcal{M} + \lambda_1)^2 \text{Tr} \left(\mathbf{G}^2 \mathbb{E} \left[\mathbf{A}_1 \mathbf{A}_1^\top \right] \right) \leq (\mathcal{M} + \lambda_1)^2 \bar{\nu} \cdot \text{Tr} \left(\mathbf{G}^2 \right). \tag{15}$$

Plugging (13), (14) and (15) into (12) tells us that

$$\begin{aligned}
& \text{Tr} \left(\mathbb{E} \left[\left(\mathbf{I} + \eta_1 \mathbf{A}_1^\top \right) \mathbf{G} \left(\mathbf{I} + \eta_1 \mathbf{A}_1 \right) \left(\mathbf{I} + \eta_1 \mathbf{A}_1^\top \right) \mathbf{G} \left(\mathbf{I} + \eta_1 \mathbf{A}_1 \right) \right] \right) \\
&\leq \text{Tr} \left(\mathbf{G}^2 \right) + 4\eta_1 \lambda_1 \text{Tr} \left(\mathbf{G}^2 \right) + 5\eta_1^2 \bar{\nu} \cdot \text{Tr} \left(\mathbf{G}^2 \right) \\
&\quad + 4\eta_1^3 (\mathcal{M} + \lambda_1) \bar{\nu} \cdot \text{Tr} \left(\mathbf{G}^2 \right) + \eta_1^4 (\mathcal{M} + \lambda_1)^2 \bar{\nu} \cdot \text{Tr} \left(\mathbf{G}^2 \right) \\
&= \left(1 + 4\eta_1 \lambda_1 + 5\eta_1^2 \bar{\nu} + 4\eta_1^3 (\mathcal{M} + \lambda_1) \bar{\nu} + \eta_1^4 (\mathcal{M} + \lambda_1)^2 \bar{\nu} \right) \text{Tr} \left(\mathbf{G}^2 \right) \\
&\leq \exp \left(4\eta_1 \lambda_1 + 10\eta_1^2 \bar{\nu} \right) \text{Tr} \left(\mathbf{G}^2 \right)
\end{aligned}$$

where in the last line we used that $\eta_i \leq \frac{1}{4\max\{\mathcal{M}, \lambda_1\}}$ and that $1 + x \leq \exp(x)$

Using the value $\mathbf{G} = \mathbf{G}_{t-1} = \mathbf{W}_{t,t-1}^\top \mathbf{v}_1 \mathbf{v}_1^\top \mathbf{W}_{t,t-1}$ and plugging the above into (11), we have

$$\begin{aligned} \gamma_t &= \text{Tr} \left(\mathbb{E} \left[\left(\mathbf{I} + \eta_1 \mathbf{A}_1^\top \right) \mathbf{G}_{t-1} \left(\mathbf{I} + \eta_1 \mathbf{A}_1 \right) \left(\mathbf{I} + \eta_1 \mathbf{A}_1^\top \right) \mathbf{G}_{t-1} \left(\mathbf{I} + \eta_1 \mathbf{A}_1 \right) \right] \right) \\ &\leq \exp(4\eta_1 \lambda_1 + 10\eta_1^2 \bar{\mathcal{V}}) \mathbb{E} [\text{Tr}(\mathbf{G}_{t-1}^2)] = \exp(4\eta_1 \lambda_1 + 10\eta_1^2 \bar{\mathcal{V}}) \gamma_{t-1}, \end{aligned}$$

where we used the fact that $\gamma_{t-1} = \mathbb{E} [\text{Tr}(\mathbf{G}_{t-1}^2)]$. Since $\gamma_0 = 1$, induction proves the lemma. \square

We now have everything to prove Theorem 3.1.

Proof of Theorem 3.1. As discussed in Section 3 the main idea of this proof is to use that Algorithm 1 is essentially one step of power method for the matrix \mathbf{B}_n and use Lemma 3.1 to bound the error. To this end, we lower and upper bound $\mathbf{v}_1^\top \mathbf{B}_n \mathbf{B}_n^\top \mathbf{v}_1$ and $\text{Tr}(\mathbf{V}_\perp^\top \mathbf{B}_n \mathbf{B}_n^\top \mathbf{V}_\perp)$, respectively.

First, using Chebyshev's inequality, we have:

$$\mathbb{P} \left[\left| \mathbf{v}_1^\top \mathbf{B}_n \mathbf{B}_n^\top \mathbf{v}_1 - \mathbb{E} [\mathbf{v}_1^\top \mathbf{B}_n \mathbf{B}_n^\top \mathbf{v}_1] \right| > \frac{1}{\sqrt{\delta}} \sqrt{\text{Var} [\mathbf{v}_1^\top \mathbf{B}_n \mathbf{B}_n^\top \mathbf{v}_1]} \right] < \delta.$$

So with probability greater than $1 - \delta$, the following holds:

$$\begin{aligned} \mathbf{v}_1^\top \mathbf{B}_n \mathbf{B}_n^\top \mathbf{v}_1 &> \mathbb{E} [\mathbf{v}_1^\top \mathbf{B}_n \mathbf{B}_n^\top \mathbf{v}_1] - \frac{1}{\sqrt{\delta}} \sqrt{\text{Var} [\mathbf{v}_1^\top \mathbf{B}_n \mathbf{B}_n^\top \mathbf{v}_1]} \\ &= \mathbb{E} [\mathbf{v}_1^\top \mathbf{B}_n \mathbf{B}_n^\top \mathbf{v}_1] - \frac{1}{\sqrt{\delta}} \sqrt{\mathbb{E} [(\mathbf{v}_1^\top \mathbf{B}_n \mathbf{B}_n^\top \mathbf{v}_1)^2] - \mathbb{E} [\mathbf{v}_1^\top \mathbf{B}_n \mathbf{B}_n^\top \mathbf{v}_1]^2} \\ &\stackrel{\zeta_1}{\geq} \exp \left(2\lambda_1 \sum_{i=1}^n \eta_i - 4\lambda_1^2 \sum_{i=1}^n \eta_i^2 \right) \times \left(1 - \frac{1}{\sqrt{\delta}} \sqrt{\exp \left(18 \sum_{i=1}^n \eta_i^2 \bar{\mathcal{V}} \right) - 1} \right) \end{aligned} \tag{16}$$

where ζ_1 follows from Lemma 5.3 and 5.4.

Furthermore, using Lemma 5.2 and Markov's inequality, we have with probability at least $1 - \delta$,

$$\text{Tr}(\mathbf{V}_\perp^\top \mathbf{B}_t \mathbf{B}_t^\top \mathbf{V}_\perp) \leq \frac{\exp \left(\sum_{i \in [n]} 2\eta_i \lambda_2 + \eta_i^2 \bar{\mathcal{V}} \right)}{\delta} \cdot \left(d + \nu \sum_{i=1}^n \eta_i^2 \exp \left(\sum_{j \in [i]} 2\eta_j (\lambda_1 - \lambda_2) \right) \right). \tag{17}$$

Consequently with probability at least $1 - 2\delta$ both (16) and (17) hold and therefore the result follows by Lemma 3.1 and choosing a δ that is smaller by a constant. \square

6 Conclusion and Future Work

This work presented a finite sample complexity and asymptotic convergence rates for the classic Oja's algorithm for top-1 component streaming PCA that match well known matrix concentration and perturbation results for computing the top eigenvector. In fact, asymptotically our bound improves upon standard matrix Bernstein bounds by a factor of $\mathcal{O}(\log d)$. Our results are tighter than existing streaming PCA results by a factor of either $\mathcal{O}(d)$ or $\mathcal{O}(1/\text{gap})$.

Our analysis relied on a novel view of the algorithm and is technically fairly simple. We hope that our analysis opens a way to make progress on the many variants of PCA that occur in both theory and practice. In particular, we believe the following directions should be of wide interest:

- **Multiple components:** Currently, our result holds only for estimating the top eigenvector of Σ . Extension of our technique to compute top- k eigenvectors is an important future direction.
- **Rayleigh quotient:** Another standard metric to measure optimality of \mathbf{w}_n is Rayleigh quotient: $\mathbf{w}_n^\top \Sigma \mathbf{w}_n$. Converting our bounds on $\sin^2(\mathbf{w}_n, \mathbf{v}_1)$ to Rayleigh quotient loses a multiplicative factor of $\mathcal{O}(1/\text{gap})$ compared to the optimal rate. A direct analysis that does not lose this factor is an interesting open problem. Results on Rayleigh quotient may also help in obtaining sample complexity guarantees that are independent of eigenvalue gap.
- **High Probability:** This work focused on obtaining tight bounds on the error. However, the dependence of our results on success probability is quite suboptimal. One way to fix this is to run many copies of the algorithm, each with say $3/4$ success probability and then output the geometric median of the solutions, which can be done in nearly linear time [17]. However, we conjecture that a tighter analysis using our techniques might directly lead to improved dependency on success probability and possibly help solve some of the other problems mentioned above.

7 Acknowledgements

Sham Kakade acknowledges funding from the Washington Research Foundation for innovation in Data-intensive Discovery.

References

- [1] I. Jolliffe, *Principal component analysis*. Wiley Online Library, 2002.
- [2] P. M. Hall, A. D. Marshall, and R. R. Martin, “Incremental eigenanalysis for classification.” in *BMVC*, vol. 98. Citeseer, 1998, pp. 286–295.
- [3] J. Weng, Y. Zhang, and W.-S. Hwang, “Candid covariance-free incremental principal component analysis,” *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 25, no. 8, pp. 1034–1040, 2003.
- [4] D. A. Ross, J. Lim, R.-S. Lin, and M.-H. Yang, “Incremental learning for robust visual tracking,” *International Journal of Computer Vision*, vol. 77, no. 1-3, pp. 125–141, 2008.
- [5] T. Krasulina, “Method of stochastic approximation in the determination of the largest eigenvalue of the mathematical expectation of random matrices,” *Automatation and Remote Control*, pp. 50–56, 1970.
- [6] E. Oja, “Simplified neuron model as a principal component analyzer,” *Journal of mathematical biology*, vol. 15, no. 3, pp. 267–273, 1982.
- [7] A. Balsubramani, S. Dasgupta, and Y. Freund, “The fast convergence of incremental pca,” in *Advances in Neural Information Processing Systems*, 2013, pp. 3174–3182.
- [8] I. Mitliagkas, C. Caramanis, and P. Jain, “Memory limited, streaming pca,” in *Advances in Neural Information Processing Systems*, 2013, pp. 2886–2894.
- [9] C. D. Sa, C. Re, and K. Olukotun, “Global convergence of stochastic gradient descent for some non-convex matrix problems,” in *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France, 6-11 July 2015*, 2015, pp. 2332–2341.
- [10] R. Vershynin, “Introduction to the non-asymptotic analysis of random matrices,” *arXiv preprint arXiv:1011.3027*, 2010.
- [11] J. A. Tropp, “User-friendly tail bounds for sums of random matrices,” *Foundations of Computational Mathematics*, vol. 12, no. 4, pp. 389–434, 2012.
- [12] I. M. Johnstone, “On the distribution of the largest eigenvalue in principal components analysis,” *Annals of statistics*, pp. 295–327, 2001.
- [13] P.-Å. Wedin, “Perturbation bounds in connection with singular value decomposition,” *BIT Numerical Mathematics*, vol. 12, no. 1, pp. 99–111, 1972.
- [14] C. Jin, S. M. Kakade, C. Musco, P. Netrapalli, and A. Sidford, “Robust shift-and-invert preconditioning: Faster and more sample efficient algorithms for eigenvector computation,” *arXiv preprint arXiv:1510.08896*, 2015.
- [15] G. H. Golub and C. F. Van Loan, *Matrix computations*. JHU Press, 2012, vol. 3.
- [16] M. Hardt and E. Price, “The noisy power method: A meta algorithm with applications,” in *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada*, 2014, pp. 2861–2869.

- [17] M. Cohen, Y. T. Lee, G. Miller, J. Pachocki, and A. Sidford, “Geometric median in nearly linear time,” *To Appear in 48th Annual Symposium on the Theory of Computing (STOC) 2016*, 2010.
- [18] K. L. Clarkson and D. P. Woodruff, “Numerical linear algebra in the streaming model,” in *Proceedings of the forty-first annual ACM symposium on Theory of computing*. ACM, 2009, pp. 205–214.
- [19] E. Liberty, “Simple and deterministic matrix sketching,” in *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 2013, pp. 581–588.
- [20] J. Nelson and H. L. Nguyễn, “Osnap: Faster numerical linear algebra algorithms via sparser subspace embeddings,” in *Foundations of Computer Science (FOCS), 2013 IEEE 54th Annual Symposium on*. IEEE, 2013, pp. 117–126.
- [21] M. B. Cohen, J. Nelson, and D. P. Woodruff, “Optimal approximate matrix product in terms of stable rank,” *arXiv preprint arXiv:1507.02268*, 2015.
- [22] M. Ghashami, E. Liberty, J. M. Phillips, and D. P. Woodruff, “Frequent directions: Simple and deterministic matrix sketching,” *arXiv preprint arXiv:1501.01711*, 2015.
- [23] C. Boutsidis, D. Garber, Z. Karnin, and E. Liberty, “Online principal components analysis,” in *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 2015, pp. 887–901.
- [24] O. Shamir, “A stochastic PCA and SVD algorithm with an exponential convergence rate,” in *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France, 6-11 July 2015*, 2015, pp. 144–152.
- [25] D. Garber and E. Hazan, “Fast and simple pca via convex optimization,” *arXiv preprint arXiv:1509.05647*, 2015.
- [26] M. K. Warmuth and D. Kuzmin, “Randomized PCA algorithms with regret bounds that are logarithmic in the dimension,” in *Advances in Neural Information Processing Systems 19, Proceedings of the Twentieth Annual Conference on Neural Information Processing Systems, Vancouver, British Columbia, Canada, December 4-7, 2006*, 2006, pp. 1481–1488.
- [27] D. Garber, E. Hazan, and T. Ma, “Online learning of eigenvectors,” in *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France, 6-11 July 2015*, 2015, pp. 560–568.