

Subspace Estimation from Incomplete Observations: A High-Dimensional Analysis

Chuang Wang, Yonina C. Eldar, *Fellow, IEEE* and Yue M. Lu, *Senior Member, IEEE*

Abstract—We present a high-dimensional analysis of three popular algorithms, namely, Oja’s method, GROUSE and PETRELS, for subspace estimation from streaming and highly incomplete observations. We show that, with proper time scaling, the time-varying principal angles between the true subspace and its estimates given by the algorithms converge weakly to deterministic processes when the ambient dimension n tends to infinity. Moreover, the limiting processes can be exactly characterized as the unique solutions of certain ordinary differential equations (ODEs). A finite sample bound is also given, showing that the rate of convergence towards such limits is $\mathcal{O}(1/\sqrt{n})$. In addition to providing asymptotically exact predictions of the dynamic performance of the algorithms, our high-dimensional analysis yields several insights, including an asymptotic equivalence between Oja’s method and GROUSE, and a precise scaling relationship linking the amount of missing data to the signal-to-noise ratio. By analyzing the solutions of the limiting ODEs, we also establish phase transition phenomena associated with the steady-state performance of these techniques.

Index Terms—Subspace tracking, streaming PCA, incomplete data, high-dimensional analysis, scaling limit

I. INTRODUCTION

Subspace estimation is a key task in many signal processing applications. Examples include source localization in array processing, system identification, network monitoring, and image sequence analysis, to name a few. The ubiquity of subspace estimation comes from the fact that a low-rank subspace model can conveniently capture the intrinsic, low-dimensional structures of many large datasets.

In this paper, we consider the problem of estimating and tracking an unknown subspace from *streaming measurements* with many *missing entries*. The streaming setting appears in applications (e.g. video surveillances) where high-dimensional data arrive sequentially over time at high rates. It is especially relevant in dynamic scenarios where the underlying subspace to be estimated can be time-varying. Missing data is also

very common in practice. Incomplete observations may result from a variety of reasons, such as the limitations of the sensing mechanisms, constraints on power consumption or communication bandwidth, or a deliberate design feature that protects privacy of individuals by removing partial records.

GROUSE [1] and PETRELS [2] as well as the classical Oja’s method [3] are three popular algorithms for solving the subspace estimation problem. They are all streaming algorithms in the sense that they provide instantaneous, *on-the-fly* updates to their subspace estimates upon the arrival of a new data point. The three differ in their update rules: Oja’s method and GROUSE perform first-order incremental gradient descent on the Euclidean space and the Grassmannian, respectively, whereas PETRELS can be interpreted as a second-order stochastic gradient descent scheme. These algorithms have been shown to be highly effective in practice, but their performance depends on careful choice of algorithmic parameters such as the step size (for GROUSE and Oja’s method) and the discount parameter (for PETRELS). Various convergence properties of these techniques have been studied in [2], [4]–[7], but a precise analysis of their performance is still an open problem. Moreover, the important question of how the signal-to-noise ratios (SNRs), the amount of missing data, and various other algorithmic parameters affect the estimation performance is not fully understood.

As the main objective of this work, we present a tractable and asymptotically exact analysis of the dynamic performance of Oja’s method, GROUSE and PETRELS in the high-dimensional regime. Our contribution is mainly threefold:

1. *Precise analysis via scaling limits.* We show in Theorem 1 and Theorem 2 that the time-varying trajectories of the estimation errors, measured in terms of the principal angles between the true underlying subspace and the estimates given by the algorithms, converge weakly to deterministic processes, as the ambient dimension $n \rightarrow \infty$. Moreover, these deterministic limits can be characterized as the unique solutions of certain ordinary differential equations (ODEs). In addition, we provide a finite-size guarantee in Theorem 3, showing that the convergence rate towards the limits is $\mathcal{O}(1/\sqrt{n})$. Numerical simulations verify the accuracy of our asymptotic predictions. The main technical tool behind our analysis is the weak convergence theory of stochastic processes (see [8]–[12] for mathematical foundations and [13]–[15] for recent applications in related estimation problems).

2. *Insights regarding the algorithms.* In addition to providing asymptotically exact predictions of the dynamic performance of the three subspace estimation algorithms, our high-dimensional analysis leads to several insights. First, the result

C. Wang is with the John A. Paulson School of Engineering and Applied Sciences, Harvard University, Cambridge, MA 02138, USA (e-mail: chuangwang@g.harvard.edu).

Y. C. Eldar is with the Department of EE, Technion, Israel Institute of Technology, Haifa, 32000, Israel (e-mail: yonina@ee.technion.ac.il).

Y. M. Lu is with the John A. Paulson School of Engineering and Applied Sciences, Harvard University, Cambridge, MA 02138, USA (e-mail: yuelu@seas.harvard.edu).

The work of C. Wang and Y. M. Lu was supported in part by the US Army Research Office under contract W911NF-16-1-0265 and in part by the US National Science Foundation under grants CCF-1319140 and CCF-1718698. The work of Y. Eldar was supported in part by the European Union’s Horizon 2020 Research and Innovation Program under Grant 646804-ERCCOG-BNYQ. Preliminary results of this work was presented at the Signal Processing with Adaptive Sparse Structured Representations (SPARS) workshop in 2017.

of Theorem 1 implies that, despite their different update rules, Oja’s methods and GROUSE are asymptotically equivalent, with both converging to the *same* deterministic process as the dimension increases. Second, the characterization given in Theorem 2 shows that PETRELS can be examined within a common framework that incorporates all three algorithms, with the difference being that PETRELS uses an adaptive scheme to adjust its effective step sizes. Third, our limiting ODEs also reveal an (asymptotically) exact scaling relationship that links the amount of missing data to the SNR. See the discussion in Section IV-A for details.

3. *Fundamental limits and phase transitions.* Analyzing the limiting ODEs also reveals phase transition phenomena associated with the steady-state performance of these algorithms. Specifically, we provide in Propositions 1 and 2 critical thresholds for setting key algorithm parameters (as a function of the SNR and the subsampling ratio), beyond which the algorithms converge to “noninformative” estimates that are no better than mere random guesses.

The rest of the paper is organized as follows. We start by presenting in Section II-A the exact problem formulation for subspace estimation with missing data. This is followed by a brief review of the three algorithms to be analyzed in this work. The main results are presented in Section III, where we show that the dynamic performance of Oja’s method, GROUSE and PETRELS can be asymptotically characterized by the solutions of certain deterministic systems of ODEs. Numerical experiments are also provided to illustrate and verify our theoretical predictions. To place our asymptotic analysis in proper context, we discuss related work in the literature in Section III-D. We consider various implications and insights drawn from our analysis in Section IV. Due to space limitation, we only present informal derivations of the limiting ODEs and proof sketches in Section V. More technical details and the proofs of all the results presented in this paper can be found in the Supplementary Materials [16].

Notation: Throughout the paper, we use \mathbf{I}_d to denote the $d \times d$ identity matrix. For any positive semidefinite matrix \mathbf{M} , its principal squared root is written as $\mathbf{M}^{-\frac{1}{2}}$. Depending on the context, $\|\cdot\|$ denotes either the ℓ_2 norm of a vector or the spectral norm of a matrix. For any $x \in \mathbb{R}$, the floor operation $\lfloor x \rfloor$ gives the largest integer that is smaller than or equal to x . Let $\{X_n\}$ be a sequence of random variables in a general probability space. $X_n \xrightarrow{\mathcal{P}} X$ means that X_n converges in probability to a random variable X , whereas $X_n \xrightarrow{\text{weakly}} X$ means that X_n converges to X weakly (*i.e.* in law). Finally, $\mathbb{1}_{\mathcal{A}}$ denotes the indicator function for an event \mathcal{A} .

II. PROBLEM FORMULATION AND OVERVIEW OF ALGORITHMS

A. Observation Model

We consider the problem of estimating a low-rank subspace using partial observations from a data stream. At any discrete-time k , suppose that a sample vector $\mathbf{s}_k \in \mathbb{R}^n$ is generated according to

$$\mathbf{s}_k = \mathbf{U}\mathbf{c}_k + \mathbf{a}_k. \quad (1)$$

Here, $\mathbf{U} \in \mathbb{R}^{n \times d}$ is an unknown deterministic matrix whose columns form an *orthonormal* basis of a d -dimensional subspace, and $\mathbf{c}_k \in \mathbb{R}^d$ is a random vector representing the expansion coefficients in that subspace. We further assume¹ that the covariance matrix of \mathbf{c}_k is diagonal:

$$\mathbf{\Lambda}^2 = \text{diag}(\lambda_1^2, \lambda_2^2, \dots, \lambda_d^2), \quad (2)$$

where $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_d$ are some strictly positive numbers. The noise in the observations is modeled by a random vector $\mathbf{a}_k \in \mathbb{R}^n$ with zero mean and covariance matrix equal to $\sigma^2 \mathbf{I}_n$. Furthermore, \mathbf{a}_k is independent of \mathbf{c}_k . Since $\{\lambda_\ell\}_{1 \leq \ell \leq d}$ in (2) indicate the “strength” of the subspace components relative to the noise \mathbf{a}_k whose variance is σ^2 , we refer to λ_ℓ/σ as the signal-to-noise ratio (SNR) for the ℓ th component of the subspace in our subsequent discussions.

We consider the missing data case, where only a subset of the entries of \mathbf{s}_k is available. This observation process can be modeled by a diagonal matrix

$$\mathbf{\Omega}_k = \text{diag}(v_{k,1}, v_{k,2}, \dots, v_{k,n}), \quad (3)$$

where $v_{k,i} = 1$ if the i th component of \mathbf{s}_k is observed, and $v_{k,i} = 0$ otherwise. Our actual observation, denoted by \mathbf{y}_k , may then be written as

$$\mathbf{y}_k = \mathbf{\Omega}_k \mathbf{s}_k. \quad (4)$$

Given a sequence of incomplete observations $\{\mathbf{y}_k, \mathbf{\Omega}_k\}_{k \geq 0}$ arriving in a stream, we aim to estimate the subspace spanned by the columns of \mathbf{U} .

B. Oja’s Method

Oja’s method [3] is a classical algorithm for estimating low-rank subspaces from streaming samples. It was originally designed for the case where the full sample vectors \mathbf{s}_k in (1) are available. Given a collection of K such sample vectors, it is natural to use the following optimization formulation to estimate the unknown subspace:

$$\hat{\mathbf{U}} = \arg \min_{\mathbf{X}^\top \mathbf{X} = \mathbf{I}_d} \sum_{k=1}^K \min_{\mathbf{w}_k} \|\mathbf{s}_k - \mathbf{X}\mathbf{w}_k\|^2 \quad (5)$$

$$= \arg \max_{\mathbf{X}^\top \mathbf{X} = \mathbf{I}_d} \sum_{k=1}^K \text{tr}(\mathbf{X}^\top \mathbf{s}_k \mathbf{s}_k^\top \mathbf{X}), \quad (6)$$

where the equivalence between (5) and (6) is established by solving the simple quadratic problem $\min_{\mathbf{w}_k} \|\mathbf{s}_k - \mathbf{X}\mathbf{w}_k\|^2$ and substituting the solution into (5).

Oja’s method is a stochastic projected-gradient algorithm for solving (6). At each step k , let \mathbf{X}_k denote the current estimate of the subspace. Then, with the arrival of a new sample vector \mathbf{s}_k , we first update \mathbf{X}_k according to

$$\widetilde{\mathbf{X}}_k = \mathbf{X}_k + \frac{\tau_k}{n} \mathbf{s}_k \mathbf{w}_k^\top, \quad (7)$$

¹The assumption that the covariance matrix is diagonal can be made without loss of generality, after a rotation of the coordinate system. To see that, suppose \mathbf{c}_k has a general covariance matrix $\mathbf{\Sigma}$, which is diagonalized as $\mathbf{\Sigma} = \mathbf{\Phi}\mathbf{\Lambda}\mathbf{\Phi}^\top$. Here, $\mathbf{\Phi}$ is an orthonormal matrix and $\mathbf{\Lambda}$ is a diagonal matrix as in (2). The generating model (1) can then be rewritten as $\mathbf{s}_k = (\mathbf{U}\mathbf{\Phi})(\mathbf{\Phi}^\top \mathbf{c}_k) + \mathbf{a}_k$. Thus, our problem is equivalent to estimating a subspace spanned by $\mathbf{U}\mathbf{\Phi}$, and $\mathbf{\Lambda}$ is the covariance matrix of the new expansion coefficient vector $\mathbf{\Phi}^\top \mathbf{c}_k$.

where $\mathbf{w}_k = \mathbf{X}_k^\top \mathbf{s}_k$ and $\{\tau_k\}$ is a sequence of positive constants that control the step-size (or learning rate) of the algorithm. We note that, up to a scaling constant, $\mathbf{s}_k \mathbf{w}_k^\top$ in (7) is exactly equal to the gradient of the objective function $\mathbf{X}^\top \mathbf{s}_k \mathbf{s}_k^\top \mathbf{X}$ in (6) due to the new sample \mathbf{s}_k . Next, to enforce the orthogonality constraint, we compute

$$\mathbf{X}_{k+1} = \widetilde{\mathbf{X}}_k (\widetilde{\mathbf{X}}_k^\top \widetilde{\mathbf{X}}_k)^{-\frac{1}{2}}, \quad (8)$$

where $(\cdot)^{-\frac{1}{2}}$ stands for the principal square root of a positive semidefinite matrix. In practice, (8) is implemented using the QR-decomposition of $\widetilde{\mathbf{X}}_k$.

To handle the case of partially-observed samples, we can modify Oja's method in two ways [17]. First, we estimate the expansion coefficients \mathbf{w}_k in (7) by solving a least squares problem that takes into account the missing data model:

$$\widehat{\mathbf{w}}_k = \arg \min_{\mathbf{w} \in \mathbb{R}^d} \|\mathbf{y}_k - \mathbf{\Omega}_k \mathbf{X}_k \mathbf{w}\|^2, \quad (9)$$

where \mathbf{y}_k is the incomplete sample vector defined in (4), $\mathbf{\Omega}_k$ is the corresponding subsampling matrix, and \mathbf{X}_k is the current estimate of the subspace. Next, we replace the missing elements in \mathbf{y}_k by the corresponding entries in $\mathbf{X}_k \widehat{\mathbf{w}}_k$. This *imputation* step leads to an estimate of the full vector:

$$\widehat{\mathbf{y}}_k = \mathbf{y}_k + (\mathbf{I}_n - \mathbf{\Omega}_k) \mathbf{X}_k \widehat{\mathbf{w}}_k. \quad (10)$$

Replacing the original vectors \mathbf{s}_k and \mathbf{w}_k in (7) by their estimated counterparts $\widehat{\mathbf{y}}_k$ and $\widehat{\mathbf{w}}_k$ leads to the modified Oja's method, a pseudocode of which is summarized in Algorithm 1.

To ensure that we have enough observed entries in \mathbf{y}_k , we first check, with the arrival of a new partially observed vector \mathbf{y}_k , whether

$$\lambda_{\min}(\mathbf{X}_k^\top \mathbf{\Omega}_k \mathbf{X}_k) > \epsilon, \quad (11)$$

where $\epsilon > 0$ is a small positive constant. If this is indeed the case, we then perform the standard update as described above; otherwise, we ignore the new sample vector and do not change the estimate in this step. Note that, under a suitable probabilistic model for the subsampling process (see assumption (A.3) in Section III-C), one can show that (11) is satisfied with high probability as long as $\epsilon < \alpha$, where α denotes the subsampling ratio defined in assumption (A.3).

C. GROUSE

Similar to Oja's method, Grassmannian Rank-One Update Subspace Estimation (GROUSE) [1] is a first-order stochastic gradient descent algorithm for solving (5). The main difference is that GROUSE solves the optimization problem on the Grassmannian, the manifold of all subspaces with a fixed rank. One advantage of this approach is that it avoids the explicit orthogonalization step in (8), allowing the algorithm to achieve lower computational complexity.

At each step, GROUSE first finds the coefficient $\widehat{\mathbf{w}}_k$ according to (9). It then computes the reconstruction error vector

$$\mathbf{r}_k = \mathbf{y}_k - \mathbf{\Omega}_k \mathbf{p}_k. \quad (12)$$

Here, \mathbf{y}_k and $\mathbf{\Omega}_k$ are defined in (1) and (3), and

$$\mathbf{p}_k = \mathbf{X}_k \widehat{\mathbf{w}}_k, \quad (13)$$

Algorithm 1 Oja's method with imputation [3], [17]

Require: An initial estimate \mathbf{X}_0 such that $\mathbf{X}_0^\top \mathbf{X}_0 = \mathbf{I}_d$, a sequence of step-size parameters $\{\tau_k\}$ and a positive constant ϵ .

```

1:  $k := 0$ 
2: repeat
3:   if  $\lambda_{\min}(\mathbf{X}_k^\top \mathbf{\Omega}_k \mathbf{X}_k) > \epsilon$  then
4:      $\widehat{\mathbf{w}}_k := \arg \min_{\mathbf{w}} \|\mathbf{y}_k - \mathbf{\Omega}_k \mathbf{X}_k \mathbf{w}\|^2$ 
5:      $\widehat{\mathbf{y}}_k := \mathbf{y}_k + (\mathbf{I}_n - \mathbf{\Omega}_k) \mathbf{X}_k \widehat{\mathbf{w}}_k$ 
6:      $\mathbf{X}_k := \mathbf{X}_k + \frac{\tau_k}{n} \widehat{\mathbf{y}}_k \widehat{\mathbf{w}}_k^\top$ 
7:      $\mathbf{X}_{k+1} := \widetilde{\mathbf{X}}_k (\widetilde{\mathbf{X}}_k^\top \widetilde{\mathbf{X}}_k)^{-\frac{1}{2}}$ 
8:   else
9:      $\mathbf{X}_{k+1} := \mathbf{X}_k$ 
10:  end if
11:   $k := k + 1$ 
12: until termination

```

Algorithm 2 GROUSE [1]

Require: An initial estimate \mathbf{X}_0 such that $\mathbf{X}_0^\top \mathbf{X}_0 = \mathbf{I}_d$, a sequence of step-size parameters $\{\tau_k\}$ and a positive constant ϵ .

```

1:  $k := 0$ 
2: repeat
3:   if  $\lambda_{\min}(\mathbf{X}_k^\top \mathbf{\Omega}_k \mathbf{X}_k) > \epsilon$  then
4:      $\widehat{\mathbf{w}}_k := \arg \min_{\mathbf{w}} \|\mathbf{y}_k - \mathbf{\Omega}_k \mathbf{X}_k \mathbf{w}\|^2$ 
5:      $\mathbf{p}_k := \mathbf{X}_k \widehat{\mathbf{w}}_k$ 
6:      $\mathbf{r}_k := \mathbf{y}_k - \mathbf{\Omega}_k \mathbf{p}_k$ 
7:      $\theta_k := \frac{\tau_k}{n} \|\mathbf{r}_k\| \cdot \|\mathbf{p}_k\|$ 
8:      $\mathbf{X}_{k+1} := \mathbf{X}_k + \left[ \frac{(\cos(\theta_k) - 1) \mathbf{p}_k}{\|\mathbf{p}_k\|} + \frac{\sin(\theta_k) \mathbf{r}_k}{\|\mathbf{r}_k\|} \right] \frac{\widehat{\mathbf{w}}_k^\top}{\|\widehat{\mathbf{w}}_k\|}$ 
9:   else
10:     $\mathbf{X}_{k+1} := \mathbf{X}_k$ 
11:  end if
12:   $k := k + 1$ 
13: until termination

```

where $\widehat{\mathbf{w}}_k$ is defined in (9). Next, it updates the current estimate \mathbf{X}_k on the Grassmannian as

$$\mathbf{X}_{k+1} = \mathbf{X}_k + \left[\frac{(\cos(\theta_k) - 1) \mathbf{p}_k}{\|\mathbf{p}_k\|} + \frac{\sin(\theta_k) \mathbf{r}_k}{\|\mathbf{r}_k\|} \right] \frac{\mathbf{w}_k^\top}{\|\mathbf{w}_k\|},$$

where

$$\theta_k = \frac{\tau_k}{n} \|\mathbf{r}_k\| \cdot \|\mathbf{p}_k\|, \quad (14)$$

and $\{\tau_k\}_k$ is a sequence of step-size parameters. The algorithm is summarized in Algorithm 2.

D. PETRELS

When there is no missing data, an alternative to Oja's method is a classical algorithm called Projection Approximation Subspace Tracking (PAST) [18]. This method estimates the underlying subspace \mathbf{U} by solving an exponentially-weighted least-squares problem

$$\mathbf{X}_{k+1} = \arg \min_{\mathbf{X} \in \mathbb{R}^{n \times d}} \sum_{k'=1}^k \gamma^{k-k'} \|\mathbf{s}_{k'} - \mathbf{X} \mathbf{w}_{k'}\|^2, \quad (15)$$

Algorithm 3 Simplified PETRELS [2]

Require: An initial estimate of the subspace \mathbf{X}_0 , $\mathbf{R}_0 = \frac{\delta}{n} \mathbf{I}_d$ for some $\delta > 0$, and positive constants γ and ϵ .

- 1: $k := 0$
- 2: **repeat**
- 3: **if** $\lambda_{\min}(\mathbf{X}_k^\top \Omega_k \mathbf{X}_k) > \epsilon$ **then**
- 4: $\hat{\mathbf{w}}_k := \arg \min_{\mathbf{w}} \|\mathbf{y}_k - \Omega_k \mathbf{X}_k \mathbf{w}\|^2$
- 5: $\mathbf{X}_{k+1} := \mathbf{X}_k + \Omega_k (\mathbf{y}_k - \mathbf{X}_k \hat{\mathbf{w}}_k) \hat{\mathbf{w}}_k^\top \mathbf{R}_k$
- 6: $\mathbf{v}_k := \gamma^{-1} \mathbf{R}_k \hat{\mathbf{w}}_k$
- 7: $\beta_k := 1 + \alpha \hat{\mathbf{w}}_k^\top \mathbf{v}_k$
- 8: $\mathbf{R}_{k+1} := \gamma^{-1} \mathbf{R}_k - \alpha \mathbf{v}_k \mathbf{v}_k^\top / \beta_k$
- 9: **else**
- 10: $\mathbf{X}_{k+1} := \mathbf{X}_k$
- 11: $\mathbf{R}_{k+1} := \mathbf{R}_k$
- 12: **end if**
- 13: $k := k + 1$
- 14: **until** termination

where $\mathbf{w}_{k'} = \mathbf{X}_{k'}^\top \mathbf{s}_{k'}$ and $\gamma \in (0, 1]$ is a discount parameter. The solution of (15) has a simple recursive update rule

$$\mathbf{X}_{k+1} = \mathbf{X}_k + (\mathbf{s}_k - \mathbf{X}_k \mathbf{w}_k) \mathbf{w}_k^\top \mathbf{R}_k \quad (16)$$

$$\mathbf{R}_{k+1} = (\gamma \mathbf{R}_k^{-1} + \mathbf{w}_k \mathbf{w}_k^\top)^{-1}. \quad (17)$$

Moreover, one can avoid the explicit calculation of the matrix inverse in (17) by using the Woodbury identity and the fact that (17) amounts to a rank-1 update.

Parallel Subspace Estimation and Tracking by Recursive Least Squares (PETRELS) [2] extends PAST to the case of partially-observed data. The main change is that it estimates the coefficient \mathbf{w}_k in (15) using (9). In addition, it provides a parallel sub-routine in its calculations so that updates to different coordinates can be computed in a fully parallel fashion. In its most general form, PETRELS needs to maintain and update a different $d \times d$ matrix \mathbf{R}_k^i for each of the n coordinates. To reduce computational complexity, a simplified version of PETRELS is provided in [2], using a common \mathbf{R}_k for all the coordinates.

In this paper, we focus on this simplified version of PETRELS, which is summarized in Algorithm 3. Note that we introduce an additional parameter α in lines 7 and 8 of the pseudocode. The simplified algorithm given in [2] corresponds to setting $\alpha = 1$. In our analysis, we set α to be equal to the subsampling ratio defined later in (27). Empirically, we find that, with this modification, the performance of the simplified algorithm matches that of the full PETRELS algorithm when the ambient dimension n is large.

Finally, we note that the computational complexity per iteration of all three algorithms is $\mathcal{O}(nd^2)$. The most expensive step is the estimation of the missing data. One main contribution of this work is to precisely predict the asymptotic ($n \rightarrow \infty$) performance of the algorithms after $k = \lfloor tn \rfloor$ iterations, for any $t > 0$.

III. MAIN RESULTS: SCALING LIMITS

In this section, we present the main results of this work—a tractable and asymptotically exact analysis of the performance of the three algorithms reviewed in Section II.

A. Performance Metric: Principal Angles

We start by defining the performance metric we will be using in our analysis. Recall the generative model defined in (1). The ground truth subspace is represented by the matrix \mathbf{U} , whose column vectors form an orthonormal basis of that subspace. For Algorithms 1, 2, and 3, the estimated subspace at the k th step is spanned by an orthogonal matrix

$$\hat{\mathbf{U}}_k \stackrel{\text{def}}{=} \mathbf{X}_k (\mathbf{X}_k^\top \mathbf{X}_k)^{-1/2}, \quad (18)$$

where \mathbf{X}_k is the k th iterand generated by the algorithms. Note that, for Oja's method and GROUSE, $\hat{\mathbf{U}}_k = \mathbf{X}_k$ as the matrix \mathbf{X}_k is already orthogonal, whereas for PETRELS, generally $\mathbf{X}_k^\top \mathbf{X}_k \neq \mathbf{I}_d$ and thus the step in (18) is necessary.

In the special case of $d = 1$ (i.e. rank-one subspaces), both \mathbf{U} and $\hat{\mathbf{U}}_k$ are unit-norm vectors. The degree to which these vectors are aligned can be measured by their *cosine similarity*, defined as $|\mathbf{U}^\top \hat{\mathbf{U}}_k|$. This concept can be naturally extended to arbitrary $d \geq 1$. In general, the closeness of two d -dimensional subspaces may be quantified by their d principal angles [19], [20]. In particular, the cosines of the principal angles are uniquely specified as the singular values of the $d \times d$ matrix

$$\mathbf{Q}_k^{(n)} \stackrel{\text{def}}{=} \mathbf{U}^\top \hat{\mathbf{U}}_k = \mathbf{U}^\top \mathbf{X}_k (\mathbf{X}_k^\top \mathbf{X}_k)^{-1/2}. \quad (19)$$

In what follows, we shall refer to $\mathbf{Q}_k^{(n)}$ as the *cosine similarity matrix*. Since we will be studying the high-dimensional limit of $\mathbf{Q}_k^{(n)}$ as the ambient dimension $n \rightarrow \infty$, we use the superscript (n) to make the dependence of $\mathbf{Q}_k^{(n)}$ on n explicit.

B. The Scaling Limits of Stochastic Processes: Main Ideas

To analyze the performance of Algorithms 1, 2, and 3, our goal boils down to tracking the evolution of the cosine similarity matrix $\mathbf{Q}_k^{(n)}$ over time. Thanks to the streaming nature of all three methods, it is easy to see that the dynamics of their estimates \mathbf{X}_k can be modeled by homogeneous Markov chains with state space in $\mathbb{R}^{n \times d}$. Thus, being a function of \mathbf{X}_k [see (19)], the dynamics of $\mathbf{Q}_k^{(n)}$ forms a hidden Markov chain. We then show that, as $n \rightarrow \infty$ and with proper time scaling, the family of stochastic processes $\{\mathbf{Q}_k^{(n)}\}$ indexed by n converges *weakly* to a deterministic function of time that is characterized as the unique solution of some ODEs. Such convergence is known in the literature as the *scaling limits* [10], [12], [15], [21] of stochastic processes. To present our results, we first consider a simple one-dimensional example that illustrates the underlying ideas behind scaling limits. Our main convergence theorems are presented in Section III-C.

Consider a 1-D stochastic process defined by a recursion

$$q_{k+1} = q_k + \frac{\tau}{n} f(q_k) + \frac{1}{n^{(1/2)+\delta}} v_k, \quad (20)$$

where $f(\cdot)$ is a Lipschitz function, τ and δ are two positive constants, v_k is a sequence of i.i.d. random variables with zero mean and unit variance, and $n > 0$ is a constant introduced to scale the step-size and the noise variance. (This particular scaling is chosen here because it mimics the actual scaling that appears in the high-dimensional dynamics of $\mathbf{Q}_k^{(n)}$ we shall study.)

When n is large, the difference between q_k and q_{k+1} is small. In other words, we will not be able to see *macroscopic* changes unless we observe the process over a large number of steps. To accelerate the time (by a factor of n), we embed $\{q_k\}$ in continuous-time by defining a piecewise-constant process

$$q^{(n)}(t) \stackrel{\text{def}}{=} q_{\lfloor nt \rfloor}, \quad (21)$$

where $\lfloor \cdot \rfloor$ is the floor function. Here, t is the rescaled (accelerated) time: within $t \in [0, 1]$, the original discrete-time process moves n steps. Due to the scaling of the noise term in (20) (with the noise variance equal to $n^{-1-2\delta}$ for some $\delta > 0$), the rescaled stochastic process $q^{(n)}(t)$ converges to a deterministic limit function as $n \rightarrow \infty$. We illustrate this convergence behavior using the following example.

Example 1: Let us consider the special case where $f(q) = -q$. We plot in Figure 1 simulation results of $q^{(n)}(t)$ for several different values of n . We see that, as n increases, the rescaled stochastic processes $q^{(n)}(t)$ converge to a limit function (the black line in the figure), which will be denoted by $q(t)$. To prove this convergence, we first expand the recursion (20) [by using the fact that $f(q) = -q$] and get

$$q_k = (1 - \frac{\tau}{n})^k q_0 + \Delta_k, \quad (22)$$

where Δ_k is a zero-mean random variable defined as

$$\Delta_k = \frac{1}{n^{(1/2)+\delta}} \sum_{i=0}^{k-1} (1 - \frac{\tau}{n})^{k-1-i} v_i.$$

Since $\{v_i\}_i$ are independent random variables with unit variance, $\mathbb{E}(\Delta_k)^2 \leq \frac{1}{n^{1+2\delta}} (1 - (1 - \frac{\tau}{n})^2)^{-1} = \mathcal{O}(n^{-2\delta})$. It then follows from (22) that, for any $t > 0$,

$$q^{(n)}(t) = q_{\lfloor nt \rfloor} \xrightarrow{\mathcal{P}} \lim_{n \rightarrow \infty} (1 - \frac{\tau}{n})^{\lfloor nt \rfloor} q_0 = q_0 e^{-\tau t}, \quad (23)$$

where $\xrightarrow{\mathcal{P}}$ stands for convergence in probability.

From (23) we can also see why $k = \lfloor nt \rfloor$ as used in (21) is the “correct” scaling under which $q^{(n)}(t)$ has a non-trivial limit as $n \rightarrow \infty$. Specifically, if the time scaling is too large, *e.g.* by choosing $k = n^b t$ with some $b > 1$, the limit on the right-hand side of (23) becomes

$$\lim_{n \rightarrow \infty} (1 - \frac{\tau}{n})^{\lfloor n^b t \rfloor} q_0 = \begin{cases} q_0, & \text{if } t = 0 \\ 0, & \text{if } t > 0. \end{cases}$$

In this case, the limiting curve $q(t)$ will abruptly drop from q_0 to 0 at $t = 0$, leaving out any details of the transition between these two extreme values. Similarly, if the time scaling is too small, *e.g.*, $k = n^{b'} t$, with some $b' < 1$, the limiting curve $q(t) \equiv q_0$ for any finite t , again revealing no information.

For general nonlinear functions $f(q)$, we can no longer directly simplify the recursion (20) as in (22). However, similar convergence behavior of $q^{(n)}(t)$ still exists. Moreover, the limit function $q(t)$ can be characterized via an ODE. To see the origin of the ODE, we note that, for any $t > 0$ and $k = \lfloor nt \rfloor$, we may rewrite (20) as

$$\frac{q^{(n)}(t + 1/n) - q^{(n)}(t)}{1/n} = \tau f[q^{(n)}(t)] + n^{(1/2)-\delta} v_k. \quad (24)$$




Figure 1. Convergence of the 1-D stochastic process $q_t^{(n)}$ described in Example 1 to its deterministic scaling limit. Here, we use $\delta = 0.25$. For each n , the associated error bars indicate ± 1 standard deviation over 1000 independent trials.

Taking the limit $n \rightarrow \infty$ on both sides of (24) and neglecting the noise term $n^{(1/2)-\delta} v_k$, we may then write—at least in a nonrigorous way—the following ODE

$$\frac{d}{dt} q(t) = \tau f[q(t)], \quad (25)$$

which always has a unique solution due to the Lipschitz property of $f(\cdot)$. For instance, the ODE associated with the linear setting in Example 1 is $\frac{d}{dt} q(t) = -\tau q(t)$, whose unique solution $q(t) = q_0 e^{-\tau t}$ is indeed the limit established in (23). A rigorous justification of the above steps can be found in the theory of weak convergence of stochastic processes (see, for example, [12], [21]).

Returning from the above detour, we recall that the central objects of our analysis are the cosine similarity matrices $\mathbf{Q}_k^{(n)}$ defined in (19). The dynamics of the cosine similarity matrix \mathbf{Q}_k is fully analogous to the situation demonstrated by the above toy example. In our work, we show that the increment $\mathbf{Q}_{k+1} - \mathbf{Q}_k$ has a mean of order $\mathcal{O}(1/n)$ and a variance of a smaller order. Thus, with time rescaling $k = \lfloor nt \rfloor$ and $n \rightarrow \infty$, the process $\mathbf{Q}(t)$ converges to the solutions of certain limiting ODEs given in the results shown in the next subsection. This phenomenon is demonstrated in Figure 2, where we plot the cosine similarity $\mathbf{Q}_{\lfloor nt \rfloor}^{(n)}$ of GROUSE at $t = 0.5$ for different values of n . In this experiment, we set $d = 1$ and thus $\mathbf{Q}_{\lfloor nt \rfloor}^{(n)}$ reduces to a scalar. The standard deviations of $\mathbf{Q}_{\lfloor nt \rfloor}^{(n)}$ over 100 independent trials, shown as error bars in Figure 2, decrease as n increases. This indicates that the performance of these stochastic algorithms can indeed be characterized by certain deterministic limits when the dimension is high.

C. The Scaling Limits of Oja’s, GROUSE and PETRELS

To study the scaling limits of the cosine similarity matrices $\mathbf{Q}_{\lfloor nt \rfloor}^{(n)}$, we embed the discrete-time process $\mathbf{Q}_k^{(n)}$ into a continuous time process $\mathbf{Q}^{(n)}(t)$ via a simple piecewise-constant interpolation:

$$\mathbf{Q}^{(n)}(t) \stackrel{\text{def}}{=} \mathbf{Q}_{\lfloor nt \rfloor}^{(n)}. \quad (26)$$




Figure 2. Convergence of the cosine similarity $Q_k^{(n)}$ associated with GROUSE at a fixed rescaled time $t = 0.5$, as we increase n from 200 to 50,000. In this experiment, $d = 1$ and thus $Q_k^{(n)}$ reduces to a scalar, denoted by $Q_k^{(n)}$. The error bars show the standard deviation of $Q_k^{(n)}$ over 100 independent trials. In each trial, we randomly generate a subspace U , the expansion coefficients $\{c_k\}$ and the noise vector $\{a_k\}$ as in (1). The red dashed lines is the limiting value predicted by our asymptotic characterization, to be given in Theorem 1.

The main objective of this work is to establish the high-dimensional limit of $Q^{(n)}(t)$ as $n \rightarrow \infty$. Our asymptotic analysis is carried out under the following technical assumptions on the generative model (1) and the observation model (3).

- (A.1) The elements of the noise vector a_k are i.i.d. random variables with zero mean, unit variance, and finite higher-order moments;
- (A.2) c_k in (1) is a d -D random vector with zero-mean and covariance matrix Λ as given in (2). Moreover, all the higher-order moments of c_k exist and are finite, and $\{c_k\}$ is independent of $\{a_k\}$;
- (A.3) We assume that $\{v_{k,i}\}$ in the observation model (3) is a collection of independent and identically distributed binary random variables such that

$$\mathbb{P}(v_{k,i} = 1) = \alpha, \quad (27)$$

for some constant $\alpha \in (0, 1)$. Throughout the paper, we refer to α as the *subsampling ratio*. We shall also assume that the algorithmic parameter ϵ used in Algorithms 1–3 satisfies the condition $\epsilon < \alpha$.

- (A.4) The target subspace matrix U and the initial guess X_0 are both *incoherent* in the sense that

$$\sum_{i=1}^n \sum_{j=1}^d U_{i,j}^4 \leq \frac{C}{n} \quad \text{and} \quad \sum_{i=1}^n \sum_{j=1}^d X_{0,i,j}^4 \leq \frac{C}{n}, \quad (28)$$

where $U_{i,j}$ and $X_{0,i,j}$ denote the (i, j) th entries of U and X_0 , respectively, and C is a generic constant that does not depend on n .

- (A.5) The initial cosine similarity $Q_0^{(n)}$ converges entrywise and in probability to a deterministic matrix $Q(0)$.
- (A.6) For Oja’s method and GROUSE, the step-size parameters $\tau_k = \tau(k/n)$, where $\tau(\cdot)$ is a deterministic function such that $\sup_{t>0} |\tau(t)| \leq C$ for a generic constant C that does not depend on n . For PETRELS, the discount factor

$$\gamma = 1 - \frac{\mu}{n}, \quad (29)$$

for some constant $\mu > 0$.

Assumption (A.4) requires some further explanation. Condition (28) essentially requires the basis matrix U and the initial guess X_0 to be generic. To see this, consider a U that is drawn uniformly at random from the Grassmannian for rank- d subspaces. Such a U can be generated as

$$U = X(X^\top X)^{-1/2}, \quad (30)$$

where X is an $n \times d$ random matrix whose entries are i.i.d. standard normal random variables. For such a generic choice of U , one can show that its entries $U_{i,j} \sim \mathcal{O}(1/\sqrt{n})$ and that (28) holds with high probability when n is large.

Theorem 1 (Oja’s method and GROUSE): Fix $T > 0$, and let $\{Q^{(n)}(t)\}_{t \in [0, T]}$ be the time-varying cosine similarity matrices associated with either Oja’s method or GROUSE over the finite interval $t \in [0, T]$. Under assumptions (A.1)–(A.6), we have

$$\{Q^{(n)}(t)\}_{t \in [0, T]} \xrightarrow{\text{weakly}} Q(t),$$

where $\xrightarrow{\text{weakly}}$ stands for weak convergence and $Q(t)$ is a deterministic matrix-valued process. Moreover, $Q(t)$ is the unique solution of the ODE

$$\frac{d}{dt} Q(t) = F(Q(t), \tau(t) I_d), \quad (31)$$

where $F : \mathbf{R}^{d \times d} \times \mathbf{R}^{d \times d} \mapsto \mathbf{R}^{d \times d}$ is a matrix-valued function defined as

$$F(Q, G) \stackrel{\text{def}}{=} [\alpha \Lambda^2 Q - \frac{\sigma^4}{2} Q G - Q(I_d + \frac{\sigma^2}{2} G) Q^\top \alpha \Lambda^2 Q] G. \quad (32)$$

Here α is the subsampling ratio, and Λ is the diagonal covariance matrix defined in (2).

In Section V, we present a (nonrigorous) derivation of the limiting ODE (31). Full technical details and a complete proof can be found in the Supplementary Materials [16]. An interesting conclusion of this theorem is that the cosine similarity matrices $Q^{(n)}(t)$ associated with Oja’s method and GROUSE converge to the same asymptotic trajectory. We will elaborate on this point in Section IV-A.

To establish the scaling limits of PETRELS, we need to introduce three auxiliary $d \times d$ matrices

$$\begin{aligned} A_k^{(n)} &\stackrel{\text{def}}{=} \frac{1}{n} R_k^{-1} \\ K_k^{(n)} &\stackrel{\text{def}}{=} U_k^\top X_k \\ W_k^{(n)} &\stackrel{\text{def}}{=} X_k^\top X_k, \end{aligned} \quad (33)$$

where the matrices R_k and X_k are those used in Algorithm 3. Then, the cosine similarity matrix can be expressed by $Q_k^{(n)} = K_k^{(n)} (W_k^{(n)})^{-\frac{1}{2}}$. Similar to (26), we embed the discrete-time processes $A_k^{(n)}$, $W_k^{(n)}$ and $W_k^{(n)}$ into continuous-time processes:

$$\begin{aligned} A^{(n)}(t) &\stackrel{\text{def}}{=} A_{[nt]}^{(n)} \\ K^{(n)}(t) &\stackrel{\text{def}}{=} K_{[nt]}^{(n)} \\ W^{(n)}(t) &\stackrel{\text{def}}{=} W_{[nt]}^{(n)}. \end{aligned} \quad (34)$$

The following theorem, whose proof can be found in the Supplementary Materials [16], characterizes the asymptotic dynamics of PETRELS.

Theorem 2 (PETRELS): For any fixed $T > 0$, let $\{(\mathbf{A}^{(n)}(t), \mathbf{K}^{(n)}(t), \mathbf{W}^{(n)}(t))\}_{t \in [0, T]}$ be the process on \mathbb{R}^{3d^2} defined in (34), and let $\{\mathbf{Q}^{(n)}(t)\}_{t \in [0, T]}$ be the time-varying cosine similarity matrices associated with PETRELS on the interval $t \in [0, T]$. Under assumptions (A.1)–(A.6), we have

$$\begin{aligned} & \{(\mathbf{A}^{(n)}(t), \mathbf{K}^{(n)}(t), \mathbf{W}^{(n)}(t))\}_{t \in [0, T]} \\ & \xrightarrow{\text{weakly}} (\mathbf{A}(t), \mathbf{K}(t), \mathbf{W}(t)), \text{ and} \\ & \{\mathbf{Q}^{(n)}(t)\}_{t \in [0, T]} \xrightarrow{\text{weakly}} \mathbf{Q}(t), \end{aligned}$$

as $n \rightarrow \infty$, where

$$\mathbf{Q}(t) = \mathbf{K}(t)\mathbf{W}(t)^{-\frac{1}{2}}, \quad (35)$$

and $(\mathbf{A}(t), \mathbf{K}(t), \mathbf{W}(t))$ is the unique solution of the following system of coupled ODEs:

$$\begin{aligned} \frac{d}{dt} \mathbf{A}(t) &= J_1(\mathbf{A}(t), \mathbf{K}(t), \mathbf{W}(t)), \\ \frac{d}{dt} \mathbf{K}(t) &= J_2(\mathbf{A}(t), \mathbf{K}(t), \mathbf{W}(t)), \\ \frac{d}{dt} \mathbf{W}(t) &= J_3(\mathbf{A}(t), \mathbf{K}(t), \mathbf{W}(t)). \end{aligned} \quad (36)$$

Here, J_1 , J_2 and J_3 are functions defined by

$$\begin{aligned} J_1(\mathbf{A}, \mathbf{K}, \mathbf{W}) &\stackrel{\text{def}}{=} \mathbf{W}^{-1}(\mathbf{K}^\top \alpha \Lambda^2 \mathbf{K} + \sigma^2 \mathbf{W}) \mathbf{W}^{-1} - \mu \mathbf{A} \\ J_2(\mathbf{A}, \mathbf{K}, \mathbf{W}) &\stackrel{\text{def}}{=} (\alpha \Lambda^2 + \sigma^2 \mathbf{I}_d) \mathbf{K} \mathbf{W}^{-1} \mathbf{A}^{-1} \\ &\quad - \mathbf{K} \mathbf{W}^{-1} (\mathbf{K}^\top \alpha \Lambda^2 \mathbf{K} + \sigma^2 \mathbf{W}) \mathbf{W}^{-1} \mathbf{A}^{-1} \\ J_3(\mathbf{A}, \mathbf{K}, \mathbf{W}) &\stackrel{\text{def}}{=} \sigma^2 \mathbf{A}^{-1} \mathbf{W}^{-1} \\ &\quad (\mathbf{K}^\top \alpha \Lambda^2 \mathbf{K} + \sigma^2 \mathbf{W}) \mathbf{W}^{-1} \mathbf{A}^{-1}, \end{aligned} \quad (37)$$

where $\mu > 0$ is the constant given in (29).

Theorem 1 and Theorem 2 establish the scaling limits of Oja's method, GROUSE and PETRELS, respectively, as $n \rightarrow \infty$. In practice, the dimension n is always finite, and thus the actual trajectories of the performance curves will fluctuate around their asymptotic limits. To bound such fluctuations via a finite-sample analysis, we first need to slightly strengthen assumption (A.5) as follows:

(A.7) Let $\mathbf{Q}_0^{(n)}$ be the initial cosine similarity matrices. There exists a fixed matrix $\mathbf{Q}(0)$ such that

$$\mathbb{E} \left\| \mathbf{Q}_0^{(n)} - \mathbf{Q}(0) \right\|_2 \leq C n^{-1/2},$$

where $\|\cdot\|_2$ denotes the spectral norm of a matrix and $C > 0$ is a constant that does not depend on n .

Theorem 3 (Finite Sample Analysis): Let $\mathbf{Q}^{(n)}(t)$ be the time-varying cosine similarity matrices associated with Oja's method, GROUSE, or PETRELS, respectively. Let $\mathbf{Q}(t)$ denote the corresponding scaling limit given in (31) and (35), respectively. Fix any $T > 0$. Under assumptions (A.1)–(A.4), (A.6)–(A.7), for any $t \in [0, T]$, we have

$$\sup_{n \geq 1} \mathbb{E} \left\| \mathbf{Q}^{(n)}(t) - \mathbf{Q}(t) \right\|_2 \leq \frac{C(T)}{\sqrt{n}}, \quad (38)$$

where $C(T)$ is a constant that can depend on the terminal time T but not on n .

The above theorem, whose proof can be found in the Supplementary Materials [16], shows that the rate of convergence towards the scaling limits is $\mathcal{O}(1/\sqrt{n})$.

Example 2: To demonstrate the accuracy of the asymptotic characterizations given in Theorem 1 and Theorem 2, we compare the actual performance of the algorithms against their theoretical predictions in Figure 3. In our experiments, we generate a random orthogonal matrix \mathbf{U} according to (30) with $n = 20,000$ and $d = 4$. For Oja's method and GROUSE, we use a constant step size $\tau = 0.5$. For PETRELS, the discount factor is $\gamma = 1 - \mu/n$ with $\mu = 5$, and $\mathbf{R}_0 = \frac{\delta}{n} \mathbf{I}_d$ with $\delta = 10$. The covariance matrix is set to

$$\mathbf{\Lambda} = \text{diag}\{5, 4, 3, 2\},$$

the subsampling ratio is $\alpha = 0.5$ and the variance of the background noise $\sigma^2 = 1$. Figure 3(a) shows the evolutions of the cosines of the 4 principal angles between \mathbf{U} and the estimates given by Oja's method (shown as crosses) and GROUSE (shown as circles). We compute the theoretical predictions of the principal angles by performing a SVD of the limiting matrices $\mathbf{Q}(t)$ as specified by the ODE (31). (In fact, this ODE has a simple analytical solution; see Section IV-B for details.) Figure 3(b) shows similar comparisons between PETRELS and its corresponding theoretical predictions. In this case, we solve the limiting ODEs (40) and (41) numerically.

D. Related Work

The problem of estimating and tracking low-rank subspaces has received a lot of attention recently in the signal processing and learning communities. Under the setting of fully observed data, an earlier work [22] studied a block-version of Oja's method and provided a sample complexity estimate for the case of $d = 1$. Similar analysis is available for general d -dimensional subspaces [23], [24]. The streaming version of Oja's method and its sample complexities have also been extensively studied, e.g., [25]–[29].

For the case of incomplete observations, the sample complexity of a block version of Oja's method with missing data is analyzed in [30] under the same generative model as in (1). In [7], the authors provide the sample complexity for learning a low-rank subspace from subsampled data under a nonparametric model: the complete data vectors are assumed to be i.i.d. samples from a general probability distribution on a subset of the Euclidean unit ball in \mathbb{R}^n . Different from that work, we study the dynamics of performance of the three algorithms rather than sample complexity. In the streaming setting, Oja's method, GROUSE, PETRELS are three popular algorithms for tackling the challenge of subspace learning with partial information. Other interesting approaches include online matrix completion methods [31]–[33]. See [17] for a recent review of relevant literature in this area. Local convergence of GROUSE is given in [4], [5]. Global convergence of GROUSE is established in [6] under the noiseless setting. In general, establishing finite sample global performance guarantees for GROUSE and other algorithms such as Oja's and PETRELS in the missing data case is still an open problem.

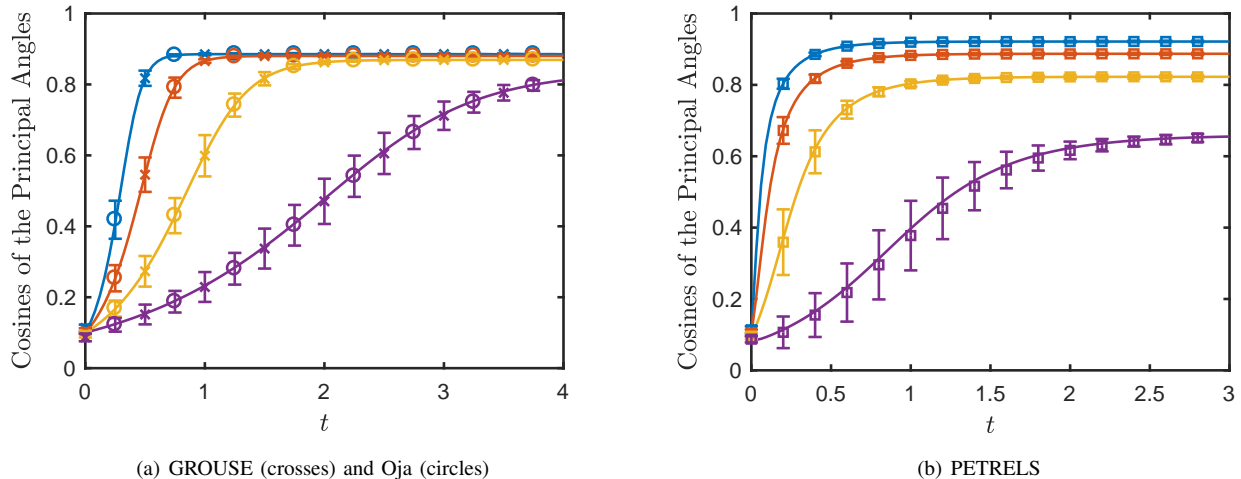


Figure 3. Numerical simulations vs. asymptotic characterizations. (a) Results for Oja’s method and GROUSE, where the solid lines are the theoretical predictions of the cosines of 4 principal angles by the solution of the ODE (31). The crosses (for Oja’s method) and circles (for GROUSE) show the simulation results averaged over 100 independent trials. In each trial, we randomly generate a subspace U as in (30), the expansion coefficients $\{c_k\}$ and the noise vector $\{a_k\}$. The error bars indicate ± 2 standard deviations. (b) Similar comparisons of numerical simulations and theoretical predictions for PETRELS.

Unlike most work in the literature that seeks to establish finite-sample performance guarantees for various subspace estimation algorithms, our results in this paper provide an asymptotically exact characterization of three popular methods in the high-dimensional limit. The main technical tool behind our analysis is the weak convergence of stochastic processes towards their scaling limits that are characterized by deterministic ODEs or stochastic differential equations (see, *e.g.*, [8]–[10], [15]).

Using ODEs to analyze stochastic recursive algorithms has a long history [34], [35]. An ODE analysis of an early subspace tracking algorithm was given in Yang [36], and this result was adapted to analyze PETRELS for the nonsampled case [2]. Our results in this paper differ from previous analysis in several important aspects. First, it handles the more challenging case of incomplete observations. Second, previous ODE analysis in Yang [36] and Chi *et al.* [2] keeps the ambient dimension n fixed and studies the asymptotic limit as the step size tends to 0. The resulting ODEs involve $\mathcal{O}(n)$ variables. In contrast, our analysis studies the limit as the dimension $n \rightarrow \infty$, and the resulting ODEs only involve at most $3d^2$ variables, where d is the dimension of the subspace which, in many practical situations, is a small constant. This low-dimensional characterization makes our limiting results more practical to use, especially when the ambient dimension n is large.

Our results build upon a previous work [13] that obtains the ODE limits of Oja’s method under the fully observed model, for the special case of $d = 1$. This paper extends that line of work to the more challenging case with missing data, and for arbitrary d . The partially observed setting presents several technical difficulties: First, unlike the fully observed case with Gaussian noise where the cosine similarity Q_k forms a low-dimensional closed Markov chain, the situation is more complicated under partial observations. Due to subsampling (which breaks rotational symmetry), the cosine similarity Q_k

no longer forms a closed Markov chain at any finite dimension n . Only when we take the asymptotic limit $n \rightarrow \infty$ do we obtain a closed dynamics given by the ODE limits. Second, in the partially observed case, an extra estimation step of the missing data should be made. This presents some additional technical challenges in the performance analysis.

It is important to point out a limitation of our asymptotic analysis: we require the initial estimate X_0 to be asymptotically correlated with the true subspace U . To see why this is an issue, we note that if the initial cosine similarity matrix $Q(0) = \mathbf{0}$ [*i.e.*, a fully uncorrelated initial estimate, which can be obtained by setting X to be an $n \times d$ i.i.d. Gaussian matrix, and assigning $X_0 = X(X^T X)^{-1/2}$], then the ODEs in Theorems 1 and 2 only provide a trivial solution $Q(t) \equiv 0$, yielding no useful information. Strictly speaking, our asymptotic predictions are still valid in this random initialization case, since Assumptions (A.4) and (A.5) are satisfied, but the predictions are not useful as $Q(t) \equiv 0$ for any finite t . It means that, starting from a random initialization, the algorithms cannot escape from the initial region in $\mathcal{O}(n)$ iterations. The utility of our results is to predict and trace the performance of the algorithms in the second phase of the dynamics, after the algorithms already have some estimates that are correlated with the true subspace. In practice, a correlated initial estimate can be obtained by performing PCA on a small batch of samples [7], [37]; it may also be available from additional side information about the true subspace U . Therefore, the requirement that $Q(0)$ be invertible is not overly restrictive. In practice, we observe in numerical simulations that, under sufficiently high SNRs, Oja’s method, GROUSE and PETRELS can successfully estimate the subspace by starting from random initial guesses that are uncorrelated with U . Extending our analysis framework to handle the case of random initial estimates is an important line of future work.

IV. IMPLICATIONS OF HIGH-DIMENSIONAL ANALYSIS

The scaling limits presented in Section III provide asymptotically exact characterizations of the dynamic performance of Oja's method, GROUSE, and PETRELS. In this section, we discuss implications of these results. Analyzing the limiting ODEs also reveals the fundamental limits and phase transition phenomena associated with the steady-state performance of these algorithms.

A. Algorithmic Insights

By examining Theorem 1 and Theorem 2, we draw the following conclusions regarding the three subspace estimation algorithms.

1. Connections and differences between the algorithms.

Theorem 1 implies that, as $n \rightarrow \infty$, Oja's method and GROUSE converge to the same deterministic limit process characterized as the solution of the ODE (31). This result is somewhat surprising, as the update rules of the two methods (see Algorithm 1 and Algorithm 2) appear to be sufficiently different.

Theorem 2 shows that PETRELS is also intricately connected to the other two algorithms. We here consider a special case that $\mathbf{K}(0) = \mathbf{U}^\top \mathbf{X}_0$ is a diagonal matrix. In this case, $\mathbf{A}(t)$, $\mathbf{K}(t)$ and $\mathbf{W}(t)$ will also be diagonal matrices for any $t > 0$. Define

$$\mathbf{G}(t) \stackrel{\text{def}}{=} \mathbf{W}^{-\frac{1}{2}} \mathbf{A}^{-1} \mathbf{W}^{-\frac{1}{2}}. \quad (39)$$

One can show that the evolution of $\mathbf{G}(t)$ and $\mathbf{Q}(t)$ as defined in (35) is also governed by a first-order ODE, which can be deduced from (36) as

$$\frac{d}{dt} \mathbf{Q}(t) = F(\mathbf{Q}(t), \mathbf{G}(t)), \quad (40)$$

$$\frac{d}{dt} \mathbf{G}(t) = H(\mathbf{Q}(t), \mathbf{G}(t)). \quad (41)$$

Here, F is the function defined in (32) and H is a function defined by

$$H(\mathbf{Q}, \mathbf{G}) \stackrel{\text{def}}{=} \mathbf{G} \left[\mu - \mathbf{G}(\sigma^2 \mathbf{G} + \mathbf{I}_d)(\mathbf{Q}^\top \alpha \mathbf{\Lambda}^2 \mathbf{Q} + \sigma^2 \mathbf{I}_d) \right]. \quad (42)$$

The ODE (40) of the cosine similarity matrix $\mathbf{Q}(t)$ for PETRELS has exactly the same form as the one for GROUSE and Oja's method shown in (31), except for the fact that the nonadaptive step-size $\tau(t) \mathbf{I}_d$ in (31) is replaced by a $d \times d$ matrix $\mathbf{G}(t)$, itself governed by an ODE (40). Thus, $\mathbf{G}(t)$ in PETRELS can be viewed as an adaptive scheme for adjusting the step-size.

To investigate how $\mathbf{G}(t)$ evolves, we run an experiment for $d = 1$. In this case, the quantities $\mathbf{Q}(t)$, $\mathbf{G}(t)$ and $\mathbf{\Lambda}$ reduce to three scalars, denoted by $Q(t)$, $G(t)$, and λ , respectively. Figure 4 shows the dynamics of PETRELS to recover this 1-D subspace. It shows that $G(t)$ increases initially, which helps boost the convergence speed. As $Q(t)$ increases (meaning the estimates becoming more accurate), however, the effective step-size $G(t)$ gradually decreases, in order to help $Q(t)$ reach a higher steady-state value.

2. *Subsampling vs. the SNR.* The ODEs in Theorems 1 and 2 also reveal an interesting (asymptotic) equivalence between




Figure 4. Monte Carlo simulations of the PETRELS algorithm v.s. asymptotic predictions obtained by the limiting ODEs given in Theorem 2 for $d = 1$. In this case, the two matrices $\mathbf{Q}(t)$ and $\mathbf{G}(t)$ reduce to two scalars $Q(t)$ and $G(t)$. The variable $G(t)$ acts as an effective step-size, which adaptively adjusts its value according to the change in $Q(t)$. The error bars shown in the figures represent one standard deviation over 50 independent trials. The signal dimension is $n = 10^4$.

the subsampling ratio α and the SNR as specified by the matrix $\mathbf{\Lambda}$. To see this, we observe from the definition of the functions F in (32) and J_1, J_2, J_3 in (37) that α always appears together with $\mathbf{\Lambda}$ in the form of a product $\alpha \mathbf{\Lambda}^2$. This implies that an observation model with subsampling ratio α and signal strength $\mathbf{\Lambda}$ will have the same asymptotic performance as a different model with subsampling ratio $\hat{\alpha}$ and signal strength $\sqrt{\alpha/\hat{\alpha}} \mathbf{\Lambda}$. In simpler terms, having missing data is asymptotically equivalent to lowering the signal strength in the fully-observable setting. We note that [7] pointed out a connection between the subsampling ratio and the sample complexity. Here, we reveal a connection between the subsampling ratio and the SNR.

B. Oja's Method and GROUSE: Analytical Solutions and Phase Transitions

Next, we investigate the dynamics of Oja's method and GROUSE by studying the solution of the ODE given in Theorem 1. To that end, we consider a change of variables by defining

$$\mathbf{P}(t) \stackrel{\text{def}}{=} [\mathbf{Q}(t) \mathbf{Q}^\top(t)]^{-1}. \quad (43)$$

One may deduce from (31) that the evolution of $\mathbf{P}(t)$ is also governed by a first-order ODE:

$$\frac{d}{dt} \mathbf{P}(t) = \mathbf{A}(t) - \mathbf{P}(t) \mathbf{B}(t) - \mathbf{B}(t) \mathbf{P}(t), \quad (44)$$

where

$$\mathbf{A}(t) = \tau(t)[2 + \tau(t)\sigma^2] \alpha \mathbf{\Lambda}^2 \quad (45)$$

$$\mathbf{B}(t) = \tau(t)(\alpha \mathbf{\Lambda}^2 - \frac{\tau(t)}{2} \sigma^4 \mathbf{I}_d) \quad (46)$$

are two diagonal matrices. Thanks to the linearity of (44), it admits an analytical solution

$$\mathbf{P}(t) = e^{-\int_0^t \mathbf{B}(r) dr} \mathbf{P}(0) e^{-\int_0^t \mathbf{B}(r) dr} + \int_0^t \mathbf{A}(s) e^{-2\int_s^t \mathbf{B}(r) dr} ds. \quad (47)$$

Note that the first two terms on the right-hand side of (47) represent the influence of the initial estimate $\mathbf{P}(0) = [\mathbf{Q}(0)\mathbf{Q}^\top(0)]^{-1}$ on the current state at t . In the special case of the algorithms using a constant step size, *i.e.*, $\tau(t) \equiv \tau > 0$, the solution (47) may be further simplified as

$$\mathbf{P}(t) = e^{-t\mathbf{B}}\mathbf{P}(0)e^{-t\mathbf{B}} + \mathbf{Z}(t), \quad (48)$$

where $\mathbf{Z}(t) = \text{diag}\{z_1(t), \dots, z_d(t)\}$ with

$$z_\ell(t) = \frac{(2 + \tau\sigma^2)\alpha\lambda_\ell^2}{2\alpha\lambda_\ell^2 - \tau\sigma^4} \left(1 - e^{\tau(2\alpha\lambda_\ell^2 - \tau\sigma^4)t}\right) \quad (49)$$

for $1 \leq \ell \leq d$. Note that if $2\alpha\lambda_\ell^2 - \tau\sigma^4 = 0$ for some ℓ , then the above expression for z_ℓ is understood via the convention that $(1 - e^{-\tau t})/0 = \tau t$.

The formula (48) reveals a phase transition phenomenon for the steady-state performance of the two algorithms as we change the step-size parameter τ . To see that, we first recall that the eigenvalues of $\mathbf{Q}^{(n)}(t)(\mathbf{Q}^{(n)}(t))^\top$ are exactly equal to the squared cosines of the d principal angles $\{\theta_\ell^n(t)\}$ between the true subspace \mathbf{U} and the estimate given by the algorithms. We say an algorithm generates an asymptotically *informative solution* if

$$\lim_{t \rightarrow \infty} \lim_{n \rightarrow \infty} \cos^2[\theta_\ell^n(t)] > 0 \quad \text{for all } 1 \leq \ell \leq d, \quad (50)$$

i.e., the steady-state estimates of the algorithms achieve nontrivial correlations with all the directions of \mathbf{U} . In contrast, a *noninformative solution* corresponds to

$$\lim_{t \rightarrow \infty} \lim_{n \rightarrow \infty} \cos^2[\theta_\ell^n(t)] = 0 \quad \text{for all } 1 \leq \ell \leq d, \quad (51)$$

in which case the steady-state estimates carry no information about \mathbf{U} . For $d > 1$, one may also have a third situation where only a subset of the directions of \mathbf{U} can be recovered (with nontrivial correlations) by the algorithm.

Proposition 1: Let $\theta_\ell^{(n)}(t)$ denotes the ℓ th principal angle between the true subspace and the estimate obtained by Oja's method or GROUSE with a constant step size τ . Under the same assumptions as in Theorem 1, we have

$$\lim_{t \rightarrow \infty} \lim_{n \rightarrow \infty} \cos^2(\theta_\ell^{(n)}(t)) = \max\left\{0, \frac{2\alpha\lambda_\ell^2 - \tau\sigma^4}{\alpha\lambda_\ell^2(2 + \tau\sigma^2)}\right\}, \quad (52)$$

where $\{\lambda_\ell\}$ are the SNR parameters defined in (2). It follows that the two algorithms provide asymptotically informative solutions *if and only if*

$$\tau < \frac{2\alpha}{\sigma^4} \min_{1 \leq \ell \leq d} \lambda_\ell^2. \quad (53)$$

Proof: Suppose the diagonal matrix \mathbf{B} in (46) has d_1 positive diagonal entries (with $0 \leq d_1 \leq d$), and $d_2 = d - d_1$ negative or zero entries. Without loss of generality, we may assume that \mathbf{B} can be split into a block form $\begin{bmatrix} \mathbf{B}_1 & \mathbf{0}_{d_1 \times d_2} \\ \mathbf{0}_{d_2 \times d_1} & -\mathbf{B}_2 \end{bmatrix}$ such that \mathbf{B}_1 only contains the positive diagonal entries, and $-\mathbf{B}_2$ only contains the nonpositive entries. Accordingly, we split the other two matrices in (48) as $\mathbf{P}(0) = \begin{bmatrix} \mathbf{P}_{1,1} & \mathbf{P}_{1,2} \\ \mathbf{P}_{2,1} & \mathbf{P}_{2,2} \end{bmatrix}$

and $\mathbf{Z}(t) = \begin{bmatrix} \mathbf{Z}_1(t) & \mathbf{0}_{d_1 \times d_2} \\ \mathbf{0}_{d_2 \times d_1} & \mathbf{Z}_2(t) \end{bmatrix}$. Applying the block matrix inverse formula to (48), we get

$$\mathbf{P}^{-1}(t) = \begin{bmatrix} \mathbf{W}_{1,1}(t) & \mathbf{W}_{1,2}(t) \\ \mathbf{W}_{2,1}(t) & \mathbf{W}_{2,2}(t) \end{bmatrix}, \quad (54)$$

where

$$\begin{aligned} \mathbf{W}_{1,1}^{-1}(t) &= e^{-t\mathbf{B}_1} \mathbf{P}_{1,1} e^{-t\mathbf{B}_1} + \mathbf{Z}_1(t) \\ &\quad - e^{-t\mathbf{B}_1} \mathbf{P}_{1,2} e^{t\mathbf{B}_2} \left(e^{t\mathbf{B}_2} \mathbf{P}_{2,2} e^{t\mathbf{B}_2} + \mathbf{Z}_2 \right)^{-1} e^{t\mathbf{B}_2} \mathbf{P}_{2,1} e^{-t\mathbf{B}_1}. \end{aligned}$$

It is easy to verify from the definitions of \mathbf{B} and \mathbf{Z} that

$$\lim_{t \rightarrow \infty} \mathbf{W}_{1,1}(t) = \text{diag} \left\{ \frac{2\alpha\lambda_1^2 - \tau\sigma^4}{\alpha\lambda_1^2(2 + \tau\sigma^2)}, \dots, \frac{2\alpha\lambda_{d_1}^2 - \tau\sigma^4}{\alpha\lambda_{d_1}^2(2 + \tau\sigma^2)} \right\}. \quad (55)$$

Similarly, we may verify that

$$\begin{aligned} \lim_{t \rightarrow \infty} \mathbf{W}_{1,2}(t) &= \mathbf{0}_{d_1 \times d_2} \\ \lim_{t \rightarrow \infty} \mathbf{W}_{2,2}(t) &= \mathbf{0}_{d_2 \times d_2}. \end{aligned} \quad (56)$$

Substituting (55) and (56) into (54) and recalling that the eigenvalues of $\mathbf{P}^{-1}(t)$ are exactly equal to the squared cosines of the principal angles, we reach (52). Applying the conditions given in (50) and (51) to (52) yields (53). \blacksquare

C. Steady-State Analysis of PETRELS

The steady-state property of PETRELS can also be obtained by studying the limiting ODEs as given in Theorem 2. The coupling of $\mathbf{Q}(t)$ and $\mathbf{G}(t)$ in (40) and (41), however, makes the analysis much more challenging. Unlike the case of Oja's method and GROUSE, we are not able to obtain closed-form analytical solutions of the ODEs for PETRELS. In what follows, we restrict our discussions to the special case of $d = 1$. This simplifies the task, as the matrix-valued ODEs (40) and (41) reduce to scalar-valued ones.

It is not hard to verify that, for any solution $\{Q(t), R(t)\}$ with an initial condition $\{Q(0), R(0)\}$, there is a symmetric solution $\{-Q(t), G(t)\}$ for the initial condition $\{-Q(0), G(0)\}$. To remove this redundancy, it is convenient to investigate the dynamics of $Q^2(t)$ and $G(t)$, which satisfy the following ODEs

$$\frac{d}{dt}[Q^2(t)] = GQ^2[2\alpha\lambda^2 - \sigma^4G - 2Q^2(\sigma^2 + \frac{1}{2}G)\alpha\lambda^2] \quad (57)$$

$$\frac{d}{dt}G(t) = G[\mu - G(G\sigma^2 + 1)(Q^2\alpha\lambda^2 + \sigma^2)]. \quad (58)$$

Figure 5 visualizes several different solution trajectories of these ODEs as the black curves in the Q - G plane. These solutions start from different initial conditions at the borders of the figures, and they converge to certain stationary points. The locations of these stationary points depend on the SNR $\{\lambda_\ell\}$, the subsampling ratio α and the discount parameter μ used by the algorithm. In Figures 5(a) and 5(b), the stationary points correspond to $Q^2 > 0$, and thus the algorithm generates asymptotically informative solutions according to the definition in (50). In contrast, Figure 5(c) and Figure 5(d) show the situations where the steady-state solutions are noninformative.

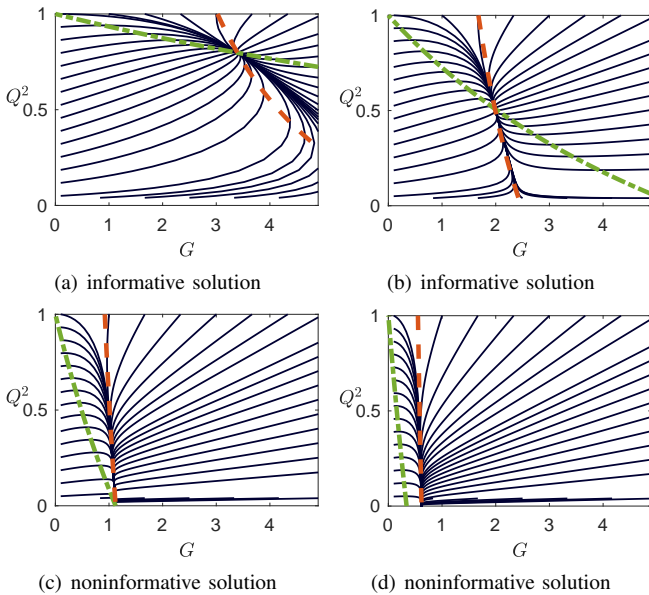


Figure 5. Phase portraits of the nonlinear ODEs in Theorem 2: The black curves are trajectories of the solutions $(Q^2(t), G(t))$ of the ODEs starting from different initial values. The green and red curves represent nontrivial solutions of the two stationary equations $\frac{dQ^2(t)}{dt} = 0$ and $\frac{dG(t)}{dt} = 0$. Their intersection point, if it exists, is a stable fixed point of the dynamical system. The fixed-points of the top two figures correspond to $Q^2(\infty) > 0$, and thus the steady-state solutions in these two cases are informative. In contrast, the fixed-points of the bottom two figures are associated with noninformative steady-state solutions with $Q^2(\infty) = 0$.

Proposition 2: Let $d = 1$. Under the same assumptions as in Theorem 2, PETRELS generates an asymptotically informative solution if and only if

$$\mu < \left(2\alpha\lambda^2/\sigma^2 + \frac{1}{2}\right)^2 - \frac{1}{4}, \quad (59)$$

where μ is the parameter defined in (29), λ denotes the SNR in (2), and α is the subsampling ratio.

Proof: It follows from Theorem 2 that verifying conditions (50) and (51) boils down to studying the fixed point of a dynamical system governed by the limiting ODEs (57) and (58). This task is in turn equivalent to setting the left-hand sides of the ODEs to zero and solving the resulting equations.

Let $\{Q^*, G^*\}$ be any solution to the equations $\frac{dQ^2}{dt} = 0$ and $\frac{dG}{dt} = 0$. From the forms of the right-hand sides of (57) and (58), we see that $\{Q^*, G^*\}$ must fall into one of the following three cases:

Case I: $G^* = 0$ and Q^* can take arbitrary values;

Case II: $Q^* = 0$ and G^* is the unique positive solution to

$$G^*(G^* + 1) = \mu; \quad (60)$$

Case III: $Q^* \neq 0$ and $G^* \neq 0$.

A local stability analysis, deferred to the end of the proof, shows that the fixed points in Case I are always unstable, in the sense that any small perturbation will make the dynamics move away from these fixed points. Thus, we just need to focus on Case II and Case III, with the former corresponding to an uninformative solution and the latter to an informative one. We will show that, under (59), a fixed point in Case III exists and it is the unique stable fixed point. That solution disappears

when (59) ceases to hold, in which case the solution in Case II becomes the unique stable fixed point.

To see why (59) provides the phase transition boundary, we note that a solution in Case III, if it exists, must satisfy $(Q^*)^2 = f(G^*)$ and $(Q^*)^2 = h(G^*)$, where

$$f(G) \stackrel{\text{def}}{=} \frac{\alpha\lambda^2 - \frac{1}{2}\sigma^4 G}{(\sigma^2 + \frac{1}{2}G)\alpha\lambda^2} \quad (61)$$

$$h(G) \stackrel{\text{def}}{=} \left(\frac{\mu}{G(\sigma^2 G + 1)} - \sigma^2\right) \frac{1}{\alpha\lambda^2}. \quad (62)$$

The above two equations are derived from $\frac{dQ^2(t)}{dt} = 0$ and $\frac{dG(t)}{dt} = 0$. In Figure 5, the functions $f(G)$ and $h(G)$ are plotted as the green and red dashed lines, respectively.

It is easy to verify from their definitions that $f(G)$ and $h(G)$ are both monotonically decreasing in the feasible region ($0 \leq Q^2 \leq 1$ and $G > 0$). Moreover, $0 = f^{-1}(1) < h^{-1}(1)$, where f^{-1} and h^{-1} denote the inverse function of f and h , respectively. Thus, a solution in Case III exists if $f^{-1}(0) > h^{-1}(0)$, which then leads to (59) after some algebraic manipulations.

Finally, we examine the local stability of the fixed points in Case I and Case II. In both cases, a fixed point (Q^*, G^*) of the 2-dimensional ODE (57) and (58) is stable if and only if

$$\frac{\partial}{\partial[Q^2]} \left[\frac{d}{dt} Q^2(t) \right] \Bigg|_{Q=Q^*, G=G^*} < 0$$

and

$$\frac{\partial}{\partial G} \left[\frac{d}{dt} G(t) \right] \Bigg|_{Q=Q^*, G=G^*} < 0,$$

where $\frac{d}{dt} Q^2(t)$ and $\frac{d}{dt} G(t)$ are the functions on the right-hand side of (57) and (58), respectively. It follows that all the Case I fixed points are always unstable, because $\frac{\partial}{\partial G} \left[\frac{d}{dt} G(t) \right] \Big|_{G=0} = \mu > 0$. Furthermore, the Case II fixed point is also unstable if (59) holds, because

$$\frac{\partial}{\partial[Q^2]} \left[\frac{d}{dt} Q^2(t) \right] \Bigg|_{Q=0, G=G^*} = 2\alpha\lambda^2 - G^* > 0,$$

where G^* is the value specified in (60). On the other hand, when (59) does not hold, the Case II fixed point becomes stable. ■

Example 3: Proposition 2 predicts a critical choice of μ (as a function of the SNR λ and the subsampling ratio α) that separates informative solutions from noninformative ones. This prediction is confirmed numerically in Figure 6. In our experiments, we set $d = 1$, $n = 10,000$. We then scan the parameter space of μ and $\alpha\lambda^2/\sigma^2$. For each choice of these two parameters on our search grid, we perform 100 independent trials, with each trial using a different realizations of \mathbf{c}_k and \mathbf{a}_k in (1) and a different \mathbf{U} drawn uniformly at random from the n -D sphere. The grayscale in Figure 6 shows the average value of the squared cosine similarity $Q(t)$ at $t = 10^3$.




Figure 6. The grayscale in the figure visualizes the steady-state squared cosine similarities of PETRELS corresponding to different values of the SNR λ^2 , the subsampling ratio α , and the step-size parameter μ . The red curve is the theoretical prediction given in Proposition 2 of a phase transition boundary, below which no informative solution can be achieved by the algorithm. The theoretical prediction matches numerical results.

V. DERIVATIONS OF THE ODES AND PROOF SKETCHES

In this section, we present a nonrigorous derivation of the limiting ODEs and sketch the main ingredients of our proofs of Theorems 1 and 2. More technical details and the complete proofs can be found in the Supplementary Materials [16].

A. Derivations of the ODE

In what follows, we show how one may derive the limiting ODE in Theorem 1. We focus on GROUSE, but the other two algorithms can be treated similarly.

Our first observation is that the dynamic of GROUSE can be modeled by a Markov chain $(\mathbf{X}_k, \mathbf{U}_k)$ on \mathbb{R}^{2dn} , where $\mathbf{U}_k \equiv \mathbf{U}$ for all k . The update rule of this Markov chain is

$$\mathbf{X}_{k+1} - \mathbf{X}_k = \left[\frac{(\cos(\theta_k) - 1)\mathbf{p}_k}{\|\mathbf{p}_k\|} + \frac{\sin(\theta_k)\mathbf{r}_k}{\|\mathbf{r}_k\|} \right] \frac{\widehat{\mathbf{w}}_k^\top}{\|\widehat{\mathbf{w}}_k\|} \mathbb{1}_{\mathcal{A}_k}, \quad (63)$$

where the vectors \mathbf{r}_k and \mathbf{p}_k are defined in (12) and (13), and

$$\begin{aligned} \mathcal{A}_k &= \{ \lambda_{\min}(\mathbf{X}_k^\top \Omega_k \mathbf{X}_k) > \epsilon \} \\ \widehat{\mathbf{w}}_k &= (\mathbf{X}_k^\top \Omega_k \mathbf{X}_k)^{-1} \mathbf{X}_k^\top \Omega_k \mathbf{s}_k. \end{aligned} \quad (64)$$

The indicator function $\mathbb{1}_{\mathcal{A}_k}$ in (63) encodes the test in line 3 of Algorithm 2. This test guarantees that the matrix inverse in (64) exists. Multiplying both sides of (63) from the left by \mathbf{U}^\top , we get

$$\mathbf{Q}_{k+1} - \mathbf{Q}_k = \frac{1}{n} \mathcal{G}_k, \quad (65)$$

where

$$\mathcal{G}_k \stackrel{\text{def}}{=} n \left[\frac{(\cos(\theta_k) - 1)\mathbf{U}^\top \mathbf{p}_k}{\|\mathbf{p}_k\|} + \frac{\sin(\theta_k)\mathbf{U}^\top \mathbf{r}_k}{\|\mathbf{r}_k\|} \right] \frac{\widehat{\mathbf{w}}_k^\top}{\|\widehat{\mathbf{w}}_k\|} \mathbb{1}_{\mathcal{A}_k} \quad (66)$$

specifies the increment of the cosine similarity from \mathbf{Q}_k to \mathbf{Q}_{k+1} .

To derive the limiting ODE, we first rewrite (65) as

$$\frac{\mathbf{Q}_{k+1} - \mathbf{Q}_k}{1/n} = \mathbb{E}_k \mathcal{G}_k + (\mathcal{G}_k - \mathbb{E}_k \mathcal{G}_k), \quad (67)$$

where \mathbb{E}_k denotes conditional expectation with respect to all the random elements encountered up to step $k - 1$, i.e.,

$\{\mathbf{c}_j, \mathbf{a}_j, \Omega_j\}_{0 \leq j \leq k-1}$ in the generative model (1) and the initial state $(\mathbf{X}_0, \mathbf{U})$. One can show that

$$\mathbb{E} \|\mathcal{G}_k - \mathbb{E}_k \mathcal{G}_k\|^2 = \mathcal{O}(1) \quad (68)$$

and

$$\mathbb{E}_k \mathcal{G}_k = F(\mathbf{Q}_k, \tau_k \mathbf{I}_d) + \mathcal{O}(1/\sqrt{n}), \quad (69)$$

where $F(\cdot, \cdot)$ is the function defined in (32), and $\|\cdot\|$ denotes the spectral norm. Substituting (69) into (67) and omitting the zero-mean difference term $(\mathcal{G}_k - \mathbb{E}_k \mathcal{G}_k)$, we get

$$\frac{\mathbf{Q}_{k+1} - \mathbf{Q}_k}{1/n} = F(\mathbf{Q}_k, \tau_k \mathbf{I}_d) + \mathcal{O}(1/\sqrt{n}). \quad (70)$$

Let $\mathbf{Q}(t)$ be a continuous-time process defined as in (26), with $t = k/n$ being the rescaled time. In an intuitive but nonrigorous way, we have $\frac{\mathbf{Q}_{k+1} - \mathbf{Q}_k}{1/n} \rightarrow \frac{d}{dt} \mathbf{Q}(t)$ as $n \rightarrow \infty$. This then gives us the ODE in (31).

In what follows, we provide some additional details behind the estimate in (69). To simplify our presentation, we first introduce a few variables. Let

$$\begin{aligned} \mathbf{Z}_k &\stackrel{\text{def}}{=} \mathbf{X}_k^\top \Omega_k \mathbf{X}_k, & \tilde{z}_k &\stackrel{\text{def}}{=} \frac{1}{n} \|\Omega_k \mathbf{s}_k\|^2, \\ \tilde{\mathbf{q}}_k &\stackrel{\text{def}}{=} \mathbf{U}^\top \Omega_k \mathbf{s}_k, & \mathbf{q}_k &\stackrel{\text{def}}{=} \mathbf{X}_k^\top \Omega_k \mathbf{s}_k, \\ \tilde{\mathbf{Q}}_k &\stackrel{\text{def}}{=} \mathbf{U}^\top \Omega_k \mathbf{X}_k, \end{aligned} \quad (71)$$

where \mathbf{s}_k and Ω_k are defined in (1) and (3). Since the columns of \mathbf{U} and \mathbf{X}_k are unit vectors, all these variables in (71) are $\mathcal{O}(1)$ quantities when $n \rightarrow \infty$. (See Lemma 5 in Supplementary Materials.)

Given its definition in (14), we rewrite θ_k used in (63) as

$$\theta_k^2 = \frac{\tau_k^2}{n} \mathbf{q}_k^\top \mathbf{Z}_k^{-2} \mathbf{q}_k \left[\tilde{z}_k - \frac{1}{n} \mathbf{q}_k^\top \mathbf{Z}_k^{-1} \mathbf{q}_k \right], \quad (72)$$

which is an $\mathcal{O}(1/n)$ quantity. Thus, it is natural to expand the two terms $\cos(\theta_k)$ and $\sin(\theta_k)$ that appear in (63) via a Taylor series expansion, which yields

$$\begin{aligned} \cos(\theta_k) &= 1 - \frac{1}{2} \theta_k^2 + \frac{1}{4!} \cos(\varphi) \theta_k^4 \\ \sin(\theta_k) &= \theta_k - \frac{1}{3!} \cos(\tilde{\varphi}) \theta_k^3. \end{aligned} \quad (73)$$

Here, the last terms in the above expansions are the remainders of the Taylor expansions, in which φ and $\tilde{\varphi}$ are some numbers between 0 and θ_k . Substituting (72) and (73) into (66) gives

$$\begin{aligned} \mathcal{G}_k &= \tau_k \left[\tilde{\mathbf{q}}_k \mathbf{q}_k^\top - (\tilde{\mathbf{Q}}_k + \frac{\tau_k}{2} \tilde{z}_k \mathbf{Q}_k) \mathbf{Z}_k^{-1} \mathbf{q}_k \mathbf{q}_k^\top \right] \mathbf{Z}_k^{-1} \mathbb{1}_{\mathcal{A}_k} \\ &\quad + \mathcal{O}(n^{-1/2}). \end{aligned} \quad (74)$$

A rigorous justification of this step is presented as Lemma 7 in the Supplementary Materials.

One can show that, as $n \rightarrow \infty$, \mathbf{Z}_k and \tilde{z}_k defined in (71) converge to $\alpha \mathbf{I}_d$ and $\alpha \sigma^2$ respectively, the quantity $\tilde{\mathbf{Q}}_k \rightarrow \alpha \mathbf{Q}_k$, and $\mathbb{1}_{\{z_k > \epsilon\}} \rightarrow 1$ (for any $\epsilon < \alpha$). Furthermore,

$$\begin{aligned} \mathbb{E} \left\| \mathbb{E}_k \mathbf{q}_k \mathbf{q}_k^\top - \alpha (\sigma^2 \mathbf{I}_d + \alpha \mathbf{Q}_k^\top \Lambda^2 \mathbf{Q}_k) \right\| &\leq C/n, \\ \mathbb{E} \left\| \mathbb{E}_k \tilde{\mathbf{q}}_k \mathbf{q}_k^\top - \alpha (\sigma^2 \mathbf{I}_d + \alpha \Lambda^2) \mathbf{Q}_k \right\| &\leq C/n, \end{aligned}$$

for some constant C , where $\|M\|$ is the spectral norm for a matrix of M . (The convergence of these variables is established in Lemma 6 in the Supplementary Materials.) Finally, by substituting the limiting values of the variables $\mathbf{Z}_k, \tilde{z}_k, \tilde{\mathbf{Q}}_k, \mathbf{q}_k \mathbf{q}_k^\top, \tilde{\mathbf{q}}_k \mathbf{q}_k^\top$ into (74), we reach the estimate in (69).

B. Main Steps of Our Proofs

The proofs of Theorems 1 and 2 follow a standard argument for proving the weak convergence of stochastic processes [8], [9], [11]. For example, to establish the scaling limit of Oja’s method and GROUSE as stated in Theorem 1, our proof consists of three main steps.

First we show that the sequence of stochastic processes $\{\{Q^{(n)}(t)\}_{0 \leq t \leq T}\}_{n=1,2,\dots}$ indexed by n is tight. The tightness property then guarantees that any such sequence must have a converging sub-sequence. Second, we prove that the limit of any converging (sub)-sequence must be a solution of the ODE (31). Third, we show that the ODE (31) admits a *unique* solution. This last property can be easily established from the fact that the function $F(\cdot, \cdot)$ on the right-hand side of (31) is a Lipschitz function (noting that $|Q_{i,j}(t)| \leq 1$ given the initial condition $|Q_{i,j}(0)| \leq 1$, where $Q_{i,j}(t)$ is the entry of $Q(t)$ at the i th row and j th column). Combining the above three steps, we may then conclude that the entire sequence of stochastic processes $\{\{Q^{(n)}(t)\}_{0 \leq t \leq T}\}_{n=1,2,\dots}$ must converge weakly to the unique solution of the ODE.

VI. CONCLUSION

In this paper, we present a high-dimensional analysis of three popular algorithms, namely, Oja’s method, GROUSE, and PETRELS, for estimating and tracking a low-rank subspace from streaming and incomplete observations. We show that, with proper time scaling, the time-varying trajectories of estimation errors of these methods converge weakly to deterministic functions of time. Such scaling limits are characterized as the unique solutions of certain ordinary differential equations (ODEs). Numerical simulations verify the accuracy of our asymptotic results. In addition to providing asymptotically exact performance predictions, our high-dimensional analysis yields several insights regarding the connections (and differences) between the three methods. Analyzing the limiting ODEs also reveals and characterizes phase transition phenomena associated with the steady-state performance of these techniques.

REFERENCES

- [1] L. Balzano, R. Nowak, and B. Recht, “Online identification and tracking of subspaces from highly incomplete information,” in *Proc. Allerton Conference on Communication, Control and Computing.*, 2010.
- [2] Y. Chi, R. Calderbank, and Y. C. Eldar, “PETRELS: Parallel subspace estimation and tracking by recursive least squares from partial observations,” *IEEE Transactions on Signal Processing*, vol. 61, no. 23, pp. 5947–5959, 2013.
- [3] E. Oja, “Simplified neuron model as a principal component analyzer,” *Journal of Mathematical Biology*, vol. 15, no. 3, pp. 267–273, 1982.
- [4] L. Balzano and S. J. Wright, “Local convergence of an algorithm for subspace identification from partial data,” *Foundations of Computational Mathematics*, vol. 15, no. 5, pp. 1279–1314, Oct. 2015.
- [5] D. Zhang and L. Balzano, “Convergence of a Grassmannian gradient descent algorithm for subspace estimation from undersampled data,” *arXiv:1610.00199*, 2016.
- [6] —, “Global Convergence of a Grassmannian Gradient Descent Algorithm for Subspace Estimation,” in *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics (AISTATS)*, vol. 51, 2016, pp. 1460–1468.
- [7] A. Gonen, D. Rosenbaum, Y. C. Eldar, and S. Shalev-Shwartz, “Subspace Learning with Partial Information,” *Journal of Machine Learning Research*, vol. 17, pp. 1–21, 2016.
- [8] S. Meleard and S. Roelly-Coppoletta, “A propagation of chaos result for a system of particles with moderate interaction,” *Stochastic Processes and their Applications*, vol. 26, pp. 317–332, Jan. 1987.
- [9] A.-S. Sznitman, “Topics in propagation of chaos,” in *Ecole d’Été de Probabilités de Saint-Flour XIX — 1989*, ser. Lecture Notes in Mathematics, P.-L. Hennequin, Ed. Springer Berlin Heidelberg, 1991, no. 1464, pp. 165–251.
- [10] S. N. Ethier and T. G. Kurtz, *Markov Processes: Characterization and Convergence*. Wiley, 1985.
- [11] P. Billingsley, *Convergence of probability measures*. John Wiley & Sons, 2013.
- [12] J. Jacod and A. Shiryaev, *Limit theorems for stochastic processes*. Springer, 2003, vol. 288.
- [13] C. Wang and Y. M. Lu, “Online learning for sparse pca in high dimensions: Exact dynamics and phase transitions,” in *Proc. IEEE Information Theory Workshop (ITW)*, Cambridge, UK, Sep. 2016.
- [14] —, “The Scaling Limit of High-Dimensional Online Independent Component Analysis,” in *Advances in Neural Information Processing Systems*, 2017.
- [15] C. Wang, J. Mattingly, and Y. M. Lu, “Scaling Limit: Exact and Tractable Analysis of Online Learning Algorithms with Applications to Regularized Regression and PCA,” *arXiv:1712.04332*, 2017.
- [16] C. Wang, Y. C. Eldar, and Y. M. Lu, “Subspace estimation from incomplete observations: A high-dimensional analysis (Main Text and Supplementary Materials),” *arXiv:1805.06834*, 2018. [Online]. Available: <https://arxiv.org/abs/1805.06834>
- [17] L. Balzano, Y. Chi, and Y. M. Lu, “A Modern Perspective on Streaming PCA and Subspace Tracking: The Missing Data Case,” *Proceedings of the IEEE*, no. to appear, 2018.
- [18] B. Yang, “Projection Approximation Subspace Tracking,” *IEEE Transactions on Signal Processing*, vol. 43, no. 1, pp. 95–107, 1995.
- [19] I. C. F. Ipsen and C. D. Meyer, “The Angle Between Complementary Subspaces,” *The American Mathematical Monthly*, vol. 102, no. 10, pp. 904–911, 1995.
- [20] F. Deutsch, “The angle between subspaces of a Hilbert space,” in *Approximation theory, wavelets and applications*. Springer, 1995, pp. 107–130.
- [21] P. Billingsley, *Convergence of probability measures*, 2nd ed., ser. Wiley series in probability and statistics. Probability and statistics section. New York: Wiley, 1999.
- [22] I. Mitliagkas, C. Caramanis, and P. Jain, “Memory limited, streaming PCA,” in *Advances in Neural Information Processing Systems*, 2013, pp. 2886–2894.
- [23] M. Hardt and E. Price, “The noisy power method: A meta algorithm with applications,” in *Advances in Neural Information Processing Systems*, 2014, pp. 2861–2869.
- [24] M.-F. Balcan, S. S. Du, Y. Wang, and A. W. Yu, “An improved gap-dependency analysis of the noisy power method,” in *Conference on Learning Theory*, 2016, pp. 284–309.
- [25] C. J. Li, M. Wang, H. Liu, and T. Zhang, “Near-optimal stochastic approximation for online principal component estimation,” *arXiv preprint arXiv:1603.05305*, 2016.
- [26] P. Jain, C. Jin, S. M. Kakade, P. Netrapalli, and A. Sidford, “Streaming pca: Matching matrix bernstein and near-optimal finite sample guarantees for oja’s algorithm,” in *Conference on Learning Theory*, 2016, pp. 1147–1164.
- [27] A. Balsubramani, S. Dasgupta, and Y. Freund, “The fast convergence of incremental pca,” in *Advances in Neural Information Processing Systems*, 2013, pp. 3174–3182.
- [28] O. Shamir, “Convergence of stochastic gradient descent for pca,” in *International Conference on Machine Learning*, 2016, pp. 257–265.
- [29] Z. Allen-Zhu and Y. Li, “First efficient convergence for streaming k-pca: a global, gap-free, and near-optimal rate,” in *Foundations of Computer Science (FOCS), 2017 IEEE 58th Annual Symposium on*. IEEE, 2017, pp. 487–492.
- [30] I. Mitliagkas, C. Caramanis, and P. Jain, “Streaming PCA with many missing entries,” *Preprint*, 2014.
- [31] A. Krishnamurthy and A. Singh, “Low-rank matrix and tensor completion via adaptive sampling,” in *Advances in Neural Information Processing Systems*, 2013, pp. 836–844.
- [32] C. Jin, S. M. Kakade, and P. Netrapalli, “Provable efficient online matrix completion via non-convex stochastic gradient descent,” in *Advances in Neural Information Processing Systems*, 2016, pp. 4520–4528.
- [33] B. Lois and N. Vaswani, “Online matrix completion and online robust pca,” in *Information Theory (ISIT), 2015 IEEE International Symposium on*. IEEE, 2015, pp. 1826–1830.

- [34] T. G. Kurtz, "Solutions of ordinary differential equations as limits of pure jump Markov processes," *Journal of Applied Probability*, vol. 7, no. 1, p. 49, Apr. 1970.
- [35] L. Ljung, "Analysis of recursive stochastic algorithms," *IEEE Transactions on Automatic Control*, vol. 22, no. 4, pp. 551–575, 1977.
- [36] B. Yang, "Asymptotic convergence analysis of the projection approximation subspace tracking algorithms," *Signal Processing*, vol. 50, no. 1–2, pp. 123–136, Apr. 1996.
- [37] E. J. Candès and B. Recht, "Exact matrix completion via convex optimization," *Foundations of Computational mathematics*, vol. 9, no. 6, p. 717, 2009.