SHIFT-INVARIANT SUBSPACE TRACKING WITH MISSING DATA

Myung Cho and Yuejie Chi

Department of Electrical and Computer Engineering
Carnegie Mellon University
Emails: {myungc,yuejiec}@andrew.cmu.edu

ABSTRACT

Subspace tracking is an important problem in signal processing that finds applications in wireless communications, video surveillance, and source localization in radar and sonar. In recent years, it is recognized that a low-dimensional subspace can be estimated and tracked reliably even when the data vectors are partially observed with many missing entries, which is greatly desirable when processing high-dimensional and high-rate data to reduce the sampling requirement. This paper is motivated by the observation that the underlying low-dimensional subspace may possess additional structural properties induced by the physical model of data, which if harnessed properly, can greatly improve subspace tracking performance. As a case study, this paper investigates the problem of tracking direction-of-arrivals from subsampled observations in a unitary linear array, where the signals lie in a subspace spanned by columns of a Vandermonde matrix. We exploit the shift-invariant structure by mapping the data vector to a latent Hankel matrix, and then perform tracking over the Hankel matrices by exploiting their low-rank properties. Numerical simulations are conducted to validate the superiority of the proposed approach over existing subspace tracking methods that do not exploit the additional shift-invariant structure in terms of tracking speed and agility.

***Index Terms*—** subspace tracking, missing data, shift-invariant subspace, Hankel matrix

1. INTRODUCTION

Subspace tracking is a classical problem in signal processing [1, 2] with applications in wireless communications, video surveillance, and source localization in radar and sonar. In a streaming setting, data vectors arrive sequentially over time, and the goal of subspace tracking is to estimate, and possibly track, a low-dimensional subspace that explains most of variability in the data without having to store all the history data in an online manner at a low complexity.

In modern high-dimensional problems, data vectors are sometimes generated at a high rate that overwhelm the processing capability of the sensing platforms, and a new challenge is that the data vectors can only be partially observed or contain many missing entries. There have been a few attempts in recent years to develop subspace tracking algorithms with missing data, such as GROUSE [3], PETRELS [4], and many others [5–7]; see [2] for an overview. These methods are made possible by the fact that the high-dimensional data vectors lie approximately in a low-dimensional subspace, and therefore can be effectively recovered even with highly incomplete observations, motivated by the success of low-rank matrix completion [8].

This paper is motivated by the observation that the low-dimensional subspace may possess additional structural properties induced by the physical model of data, which if harnessed properly, can greatly improve subspace tracking performance. With the additional structure, the number of parameters that we need to specify the low-dimensional subspace is further reduced and therefore, it is possible to improve the tracking performance with fewer observations if the structure is respected within the algorithm design.

As a case study, in array signal processing, the data snapshots collected from a uniform linear array may be modeled as a sum of a small number of complex sinusoids, which lie in a shift-invariant subspace spanned by the columns of a Vandermonde matrix. It is therefore interesting to see if this additional structure can be leveraged. In this paper, we study the problem of tracking direction-of-arrivals in a uniform linear array with randomly subsampled measurements in a streaming setting. Similar problems have been considered in the batch setting [9], however they cannot be used in the streaming setting to track changes in the scene.

Instead of directly looking at the data snapshot as a vector, we resort to the Hankel matrix constructed by the data snapshot as the first column and the last row, by folding it at certain entry. It is well-known that the resulting Hankel matrix is low-rank if the number of sinusoids is small [10], and it is possible to recover the data snapshot even if it is highly subsampled by performing low-rank matrix completion over the Hankel form [11–13]. Therefore, instead of performing subspace tracking directly over the data snapshots, we perform subspace tracking over the Hankel matrices spanned by the data snapshots, by extending the recursive least-squares formulation in the PETRELS algorithm [4] with a few important modifications. The final formulation becomes tracking a dynamic Hankel-structured tensor with low CP-rank, which is related to the algorithms for tracking unstructured tensor data in [5, 14]; yet we emphasize our focus here is to demonstrate the potential performance gain of augmenting matrix-valued data into tensor-valued data by leveraging known structures, which is not considered before.

The rest of this paper is organized as follows. Section 2 provides the problem formulation. Sections 3 presents the developed algorithm dubbed Shift-Invariant Subspace Tracking (SIST). Section 4 presents numerical experiments to validate the superior tracking performance of the proposed algorithm in direction-of-arrival estimation. Finally, we conclude in Section 5.

2. PROBLEM FORMULATION AND BACKGROUNDS

The signals impinging on a uniform linear array of length n can be modeled as a sum of r complex sinusoids, where r is the number of direction-of-arrivals. At the t-th snapshot, the signal can be written

This work is supported in part by the grants AFOSR FA9550-15-1-0205, ONR N00014-18-1-2142, NSF CIF-1826519 and ECCS-1818571.

as

$$\boldsymbol{x}[t]^\star = \sum_{i=1}^{r} c[t]_i \boldsymbol{a}(f_i) := \boldsymbol{V}\boldsymbol{c}[t], \tag{1}$$

where the atom $\boldsymbol{a}(f)$ is stated as follows:

$$\boldsymbol{a}(f) = \frac{1}{\sqrt{n}}[1,\ e^{j2\pi f}, ...,\ e^{j2\pi f(n-1)}]^T, \quad f \in [0,1). \tag{2}$$

The subspace matrix $\boldsymbol{V}$ is given as $\boldsymbol{V} = [\boldsymbol{a}(f_1), ..., \boldsymbol{a}(f_r)] \in \mathbb{C}^{n\times r}$, $f_i's \in [0,1)$ are distinct, and $\boldsymbol{c}[t] = [c[t]_i, \ldots, c[t]_r]^T \in \mathbb{C}^r$ represents the coefficient vector at time t. We assume each snapshot $\boldsymbol{x}[t]^\star$ is partially observed, where its observation pattern is indicated by the vector $\Omega[t] \in \{0,1\}^n$, where $\Omega[t]_i = 1$ if the i-th entry of $\boldsymbol{x}[t]^\star$ is observed and $\Omega[t]_i = 0$ vice versa. Our goal is to recover, and possibly track if it is changing, the underlying subspace and the corresponding frequencies from the partial observations $\{\Omega[t] \odot \boldsymbol{x}[t]^\star\}_{t=1}^{\infty}$ in a streaming fashion, where $\odot$ denotes the point-wise product.

2.1. Hankel matrix enhancement

Before continuing, we introduce the Hankel matrix enhancement for a given signal $\boldsymbol{x}^\star = [x_0^\star, x_1^\star, ..., x_{n_1+n_2-2}^\star]^T$ of length $n = n_1 + n_2 - 1$. Define $\mathcal{H}\boldsymbol{x}^\star$: $\mathbb{C}^{n_1+n_2-1} \mapsto \mathbb{C}^{n_1\times n_2}$ as the Hankel matrix obtained from $\boldsymbol{x}^\star$ whose first column and last row are filled by the entries of $\boldsymbol{x}^\star$, i.e.

$$\mathcal{H}\boldsymbol{x}^\star = \begin{bmatrix} x_0^\star & x_1^\star & x_2^\star & \ldots & x_{n_2-1}^\star \\ x_1^\star & x_2^\star & x_3^\star & \ldots & x_{n_2}^\star \\ \vdots & \vdots & \vdots & \ldots & \vdots \\ x_{n_1-1}^\star & x_{n_1}^\star & x_{n_1+1}^\star & \cdots & x_{n_1+n_2-2}^\star \end{bmatrix}. \tag{3}$$

From (1), $\mathcal{H}\boldsymbol{x}[t]^\star$ can be expressed as follows:

$$\mathcal{H}\boldsymbol{x}[t]^\star = \boldsymbol{V}_{n_1\times r}\mathrm{diag}(\{c[t]_i\}_{i=1}^r)\boldsymbol{V}_{n_2\times r}^T, \tag{4}$$

where $\boldsymbol{V}_{n_1\times r}$ and $\boldsymbol{V}_{n_2\times r}$ are the partial matrices of $\boldsymbol{V}$ taking the first n_1 and n_2 rows respectively and $\mathrm{diag}(\{c[t]_i\}_{i=1}^r)$ is the diagonal matrix whose i-th diagonal entry is $c[t]_i$. From this decomposition, we know that $\mathcal{H}\boldsymbol{x}[t]^\star$ is a rank-r matrix. With partial observations of the signal $\boldsymbol{x}[t]^\star$, it is suggested in [11, 13] that the missing entries of $\boldsymbol{x}[t]^\star$ can be perfectly estimated by applying convex and nonconvex optimization to complete the partially-observed Hankel low-rank matrix.

2.2. Shift-Invariant Subspace Tracking

In this paper, we propose to track the subspace using up to L snapshots of partial observations by the following optimization problem with respect to $\boldsymbol{A} \in \mathbb{C}^{n_1\times r}$, $\boldsymbol{B} \in \mathbb{C}^{n_2\times r}$, and $\boldsymbol{C} \in \mathbb{C}^{r\times L}$:

$$\begin{aligned} \min_{\boldsymbol{A},\boldsymbol{B},\boldsymbol{C}} \ & \frac{1}{2}\sum_{\tau=1}^{L}\lambda^{L-\tau}\left(\left\|\boldsymbol{\Omega}[\tau]\odot\left(\mathcal{H}\boldsymbol{x}[\tau]^\star - \boldsymbol{A}\,\mathrm{diag}(\boldsymbol{c}[\tau])\boldsymbol{B}^T\right)\right\|_F^2 \right.\\ & \left. + \mu_1\|\boldsymbol{c}[\tau]\|_2^2\right) + \frac{\mu_2}{2}\left(\|\boldsymbol{A}\|_F^2 + \|\boldsymbol{B}\|_F^2\right), \\ \text{s.t.} \quad & \boldsymbol{C} = [\boldsymbol{c}[1], \boldsymbol{c}[2], ..., \boldsymbol{c}[L]]. \end{aligned} \tag{5}$$

Here, λ is called the forgetting parameter, with $0 \le \lambda < 1$, and $\mu_1 > 0$, $\mu_2 > 0$ are regularization parameters. The matrix $\boldsymbol{\Omega}[\tau] \in \mathbb{R}^{n_1\times n_2}$ is an indicator matrix introduced for the projection onto the sample space in the Hankel matrix form at time τ; namely, if the l-th entry of $\boldsymbol{x}[\tau]^\star$ is observed, then, the l-th skew diagonal of $\boldsymbol{\Omega}[\tau]$ has all the same weight values. It is noted that the missing data information at time τ is incorporated by $\boldsymbol{\Omega}[\tau] \in \mathbb{R}^{n_1\times n_2}$ in (5). We call (5) as Shift-Invariant Subspace Tracking (SIST) problem using the Hankel matrix enhancement in streaming data setting. Similar formulations without the Hankel structures to (5) are introduced in [5, 14] for tracking tensor low-rank data. The major difference between the previous research in [5, 14] and ours is that we consider the structural information, i.e., Hankel structures, and seek to understand its benefits in subspace tracking performance.

3. PROPOSED ALGORITHM

In order to solve SIST problem formulated in (5), we introduce the alternating least-squares algorithm following the updating steps in [4, 14]. The overall algorithm is stated in Algorithm 1. For updating $\boldsymbol{c}[t]$, we calculate the following optimization problem with fixed $\boldsymbol{A}[t-1]$ and $\boldsymbol{B}[t-1]$:

$$\begin{aligned} \boldsymbol{c}[t] = & \operatorname*{argmin}_{\boldsymbol{c}\in\mathbb{C}^r} \ \frac{1}{2}\left\|\boldsymbol{\Omega}[t]\odot\left(\mathcal{H}\boldsymbol{x}[t]^\star - \boldsymbol{A}[t-1]\,\mathrm{diag}(\boldsymbol{c})\boldsymbol{B}[t-1]^T\right)\right\|_F^2 \\ & + \mu_1\|\boldsymbol{c}\|_2^2 \qquad (6)\\ = & \left[\sum_{(l,w)\in\boldsymbol{\Omega}[t]}\overline{\boldsymbol{g}_{l,w}[t]}\boldsymbol{g}_{l,w}[t]^T + \mu_1\boldsymbol{I}\right]^{-1}\left[\sum_{(l,w)\in\boldsymbol{\Omega}[t]}\mathcal{H}\boldsymbol{x}[t]_{l,w}^\star\overline{\boldsymbol{g}_{l,w}[t]}\right], \end{aligned}$$

where $\boldsymbol{g}_{l,w}[t] := \boldsymbol{a}^l[t-1]\odot\boldsymbol{b}_w[t-1] \in \mathbb{C}^{r\times 1}$, $\boldsymbol{a}^l[t-1] \in \mathbb{C}^{r\times 1}$ is the l-th row of $\boldsymbol{A}[t-1]$, $\boldsymbol{b}_w[t-1]$ is the w-th column of $\boldsymbol{B}[t-1]^T$. Here, the notation $\overline{\boldsymbol{g}}$ represents the complex conjugate of a vector $\boldsymbol{g}$.

After updating $\boldsymbol{c}[t]$, we update the subspace $\boldsymbol{A}[t]$ and $\boldsymbol{B}[t]$ based on the Recursive Least-Squares (RLS) algorithm, which is the technique used in [4, 14]. Updating $\boldsymbol{A}$ is conducted with fixed $\boldsymbol{c}[t]$ and $\boldsymbol{B}[t-1]$. Similarly, $\boldsymbol{B}$ is updated with fixed $\boldsymbol{c}[t]$ and $\boldsymbol{A}[t-1]$.

For $\boldsymbol{A}[t]$, we solve the following optimization problem with fixed $\boldsymbol{c}[t]$ and $\boldsymbol{B}[t-1]$:

$$\begin{aligned} & \operatorname*{minimize}_{\boldsymbol{A}\in\mathbb{C}^{n_1\times r}} \ \frac{\mu_2}{2}\|\boldsymbol{A}\|_F^2 \qquad (7)\\ & + \frac{1}{2}\sum_{\tau=1}^{t}\left(\lambda^{t-\tau}\left\|\boldsymbol{\Omega}[\tau]\odot(\mathcal{H}\boldsymbol{x}[\tau]^\star - \boldsymbol{A}\,\mathrm{diag}(\boldsymbol{c}[\tau])\boldsymbol{B}[\tau-1]^T)\right\|_F^2\right). \end{aligned}$$

Since the objective function can be decomposed over each row vector of $\boldsymbol{A}$, let us consider the optimization problem with the l-th row of $\boldsymbol{A}$, denoted by $\boldsymbol{a}^l \in \mathbb{C}^{r\times 1}$. By denoting $\mathrm{diag}(\boldsymbol{c}[\tau])\boldsymbol{b}^w[\tau-1] = \boldsymbol{q}_w[\tau] \in \mathbb{C}^{r\times 1}$, where $\boldsymbol{b}^w[\tau-1] \in \mathbb{C}^{r\times 1}$ is the w-th row of $\boldsymbol{B}[\tau-1]$, we have

$$\begin{aligned} \boldsymbol{a}^l[t] = & \operatorname*{argmin}_{\boldsymbol{a}\in\mathbb{C}^{r\times 1}} \ \frac{1}{2}\sum_{\tau=1}^{t}\sum_{w\in\boldsymbol{\Omega}[\tau]_{l,:}}\lambda^{t-\tau}\left(\mathcal{H}\boldsymbol{x}[\tau]_{l,w}^\star - \boldsymbol{a}^T\boldsymbol{q}_w[\tau]\right)^2 \\ & + \frac{\mu_2}{2}\|\boldsymbol{a}\|_2^2. \end{aligned} \tag{8}$$

where $\boldsymbol{\Omega}[\tau]_{l,:}$ represents the index set for column that has one in the l-th row of $\boldsymbol{\Omega}[\tau]$. By applying the first-order optimality condition at $\boldsymbol{a}^l[t]$, we have

$$\begin{aligned} & \left[\sum_{\tau=1}^{t}\left(\sum_{w\in\boldsymbol{\Omega}[\tau]_{l,:}}\lambda^{t-\tau}\overline{\boldsymbol{g}_w[\tau]}\boldsymbol{q}_w[\tau]^T\right) + \mu_2\boldsymbol{I}\right]\boldsymbol{a}^l[t] \\ & = \sum_{\tau=1}^{t}\sum_{w\in\boldsymbol{\Omega}[\tau]_{l,:}}\lambda^{t-\tau}\mathcal{H}\boldsymbol{x}[\tau]_{l,w}^\star\overline{\boldsymbol{g}_w[\tau]}. \end{aligned} \tag{9}$$

Algorithm 1: Shift-Invariant Subspace Tracking (SIST)

Input: λ, μ_1, μ_2, and $\mathcal{P}_\Omega \boldsymbol{x}^\star[t]$, $t = 1, ..., L$
Output: Frequency support $\boldsymbol{f}$
Initialize: $\boldsymbol{R}_l[0]^{-1} \leftarrow \boldsymbol{I}$, $\boldsymbol{Q}_w[0]^{-1} \leftarrow \boldsymbol{I}$, and random $\boldsymbol{A}[0]$, $\boldsymbol{B}[0]$, $\boldsymbol{C}[0]$
for $t = 1$ **to** L **do**
    Update $\boldsymbol{c}[t]$ via (6)
    for $l = 1$ **to** n_1 **do**
        Update $\boldsymbol{R}_l[t]$ and $\boldsymbol{a}^l[t]$ via (13) and (15)
    end
    for $w = 1$ **to** n_2 **do**
        Update $\boldsymbol{Q}_w[t]$ and $\boldsymbol{b}^w[t]$ via (17) and (16)
    end
    Extract frequencies from the subspace $\boldsymbol{A}[t]$ or $\boldsymbol{B}[t]$
end

Defining

$$\boldsymbol{R}_l[t] := \sum_{\tau=1}^{t}\left(\sum_{w\in\boldsymbol{\Omega}[\tau]_{l,:}} \lambda^{t-\tau}\overline{\boldsymbol{g}_w[\tau]}\boldsymbol{q}_w[\tau]^T\right) + \mu_2 \boldsymbol{I}, \qquad (10)$$

$$\boldsymbol{s}_l[t] := \sum_{\tau=1}^{t}\sum_{w\in\boldsymbol{\Omega}[\tau]_{l,:}} \lambda^{t-\tau}\mathcal{H}\boldsymbol{x}[\tau]^\star_{l,w}\overline{\boldsymbol{g}_w[\tau]}, \qquad (11)$$

for (9), we have

$$\boldsymbol{R}_l[t]\boldsymbol{a}^l[t] = \boldsymbol{s}_l[t]. \qquad (12)$$

Note that both $\boldsymbol{R}[t]$ and $\boldsymbol{s}_l[t]$ can be updated recursively:

$$\boldsymbol{R}_l[t] = \lambda\boldsymbol{R}_l[t-1] + \sum_{w\in\boldsymbol{\Omega}[t]_{l,:}} \overline{\boldsymbol{g}_w[t]}\boldsymbol{q}_w[t]^T + \mu_2(1-\lambda)\boldsymbol{I} \qquad (13)$$

$$\boldsymbol{s}_l[t] = \lambda\boldsymbol{s}_l[t-1] + \sum_{w\in\boldsymbol{\Omega}[t]_{l,:}} \mathcal{H}\boldsymbol{x}[t]^\star_{l,w}\overline{\boldsymbol{g}_w[t]}. \qquad (14)$$

It is easy to derive that updating $\boldsymbol{a}^l[t]$ can be conducted as

$$\begin{aligned}\boldsymbol{a}^l[t] = {} & \boldsymbol{a}^l[t-1] - \mu_2(1-\lambda)\boldsymbol{R}_l[t]^{-1}\boldsymbol{a}^l[t-1] \qquad (15)\\ & + \sum_{w\in\boldsymbol{\Omega}[t]_{l,:}}\left(\mathcal{H}\boldsymbol{x}[t]^\star_{l,w} - \boldsymbol{q}_w[t]^T\boldsymbol{a}^l[t-1]\right)\boldsymbol{R}_l[t]^{-1}\overline{\boldsymbol{g}_w[t]}.\end{aligned}$$

Then, similarly to $\boldsymbol{a}^l[t]$, for updating the w-th row vector of $\boldsymbol{B}[t]$, denoted by $\boldsymbol{b}^w[t]$, we have

$$\begin{aligned}\boldsymbol{b}^w[t] = {} & \boldsymbol{b}^w[t-1] - \mu_2(1-\lambda)\boldsymbol{Q}_w[t]^{-1}\boldsymbol{b}^w[t-1] \qquad (16)\\ & + \sum_{l\in\boldsymbol{\Omega}[t]_{:,w}}\left(\mathcal{H}\boldsymbol{x}[t]^\star_{l,w} - \boldsymbol{\beta}^l[t]^T\boldsymbol{b}^w[t-1]\right)\boldsymbol{Q}_w[t]^{-1}\overline{\boldsymbol{\beta}^l[t]},\end{aligned}$$

where $\boldsymbol{\beta}^l[\tau] \in \mathbb{C}^{1\times r} := (\boldsymbol{a}^l[\tau]^T \operatorname{diag}(\boldsymbol{c}[\tau]))^T$, $\boldsymbol{\Omega}[t]_{:,w}$ is introduced for the index set of rows that have ones in the w-th column of $\boldsymbol{\Omega}[t]$, and

$$\boldsymbol{Q}_w[t] = \lambda\boldsymbol{Q}_w[t-1] + \sum_{l\in\boldsymbol{\Omega}[t]_{:,w}} \overline{\boldsymbol{\beta}^l[t]}\boldsymbol{\beta}^l[t]^T + \mu_2(1-\lambda)\boldsymbol{I}. \qquad (17)$$

After updating $\boldsymbol{c}[t]$, $\boldsymbol{A}[t]$, and $\boldsymbol{B}[t]$, we can extract the location of frequencies from well-known frequency extraction algorithm called ESPRIT [15] by using the subspace matrix $\boldsymbol{A}[t]$ or $\boldsymbol{B}[t]$.

4. NUMERICAL EXPERIMENTS

We compare the proposed SIST algorithm to PETRELS [4] and GROUSE [3] in the direction-of-arrival estimation problem where the underlying modes are generated in two scenarios: (1) abrupt changes and (2) fast yet smooth changes. In both scenarios, we set the forgetting parameter λ for PETRELS and SIST to 0.98 and 0.1 respectively. Additionally, the parameters μ_1 and μ_2 in SIST are set to 10^{-10} for both scenarios. These choices are tuned carefully to provide desirable performance.

For the first scenario, we follow the numerical experiment setting in [4, Section VI.B]. Fig. 1 (a) shows the ground truth mode locations. For this experiment, we randomly observe 15% of the signal dimension $n = 256$ at each snapshot. The total number of snapshots L is 4000, and after every 1000 snapshots, one or more signal sources suddenly change the locations, disappear from or enter the scene. The magnitude of each signal source is chosen between 0 and 1. For measurements, we added random noise following $\mathcal{CN}(0, 0.01)$. We set a postulated rank r to 10, which is twice larger than the true rank, for all algorithms. Fig. 1 (b), (c) and (d) show the estimated mode locations with respect to the streaming index of GROUSE, PETRELS and SIST respectively. As shown in Fig. 1, both PETRELS and SIST track the subspace, yet SIST is much faster than PETRELS in terms of tracking speed.

For the second scenario, where the mode locations change smoothly yet fast, we generate three moving targets as shown in Fig. 2 (a). The ground truth signal has the modes $\boldsymbol{f} = [1 - 1/1000t, 1/1000t, 0.5 + \sin(0.01t)/2]$ with the amplitude $[0.3, 0.5, 1]$ and 1000 numbers of snapshots. Therefore, at each time t, the locations of modes are smoothly changing. For this experiment, we randomly observe 15% of the signal dimension $n = 256$ at each snapshot and added random noise to the measurements following $\mathcal{CN}(0, 0.01)$. We set the postulated rank r to 10. Fig. 2 shows the estimated mode locations with respect to the stream index. In this scenario, GROUSE and PETRELS shown in Fig. 2 (b) and (c) completely fail, while SIST shown in Fig. 2 (d) successfully tracks the modes, significantly outperforming PETRELS and GROUSE in terms of accuracy and speed.

5. CONCLUSIONS

In this paper, we study the subspace tracking problem using uniform linear array in streaming data with partial observations. In order to achieve the enhanced subspace tracking performance, we propose the Shift-Invariant Subspace Tracking (SIST) algorithm based on the Hankel matrix enhancement which is solved using the recursive least-squares method. Our algorithm is evaluated in the direction-of-arrival estimation problem, where it is demonstrated that our proposed approach outperforms existing methods including GROUSE [3] and PETRELS [4] with better agility and tracking speed. In conclusion, it is demonstrated that the performance of subspace tracking algorithms can be significantly improved by properly incorporated structural information of the subspace.

6. REFERENCES

[1] P. Comon and G. H. Golub, "Tracking a few extreme singular values and vectors in signal processing," *Proceedings of the IEEE*, vol. 78, no. 8, pp. 1327–1343, 1990.

[2] L. Balzano, Y. Chi, and Y. M. Lu, "Streaming PCA and sub-

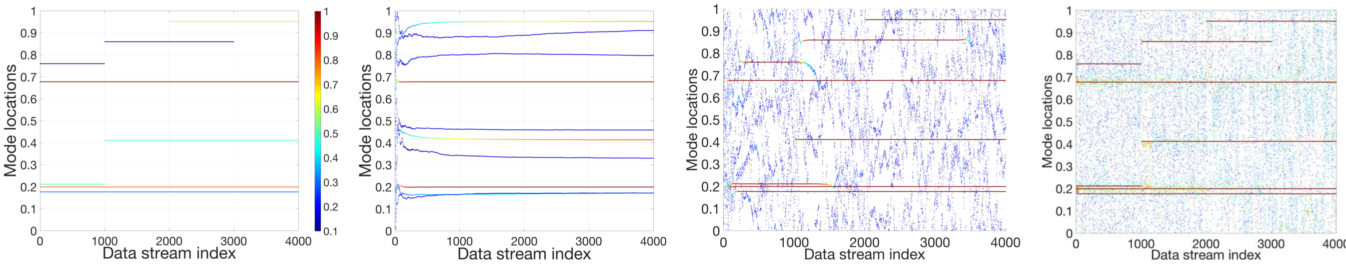

(a) Ground truth (b) GROUSE (c) PETRELS (d) SIST

Fig. 1. Tracking of mode changes in direction-of-arrival estimation under abrupt changes, using GROUSE, PETRELS and SIST: the true mode locations at each stream index is shown in (a), the estimated mode locations at each stream index for 10 modes are shown in (b), (c) and (d) for GROUSE, PETRELS and SIST respectively. All changes are identified and tracked successfully by PETRELS and SIST, but not by GROUSE.

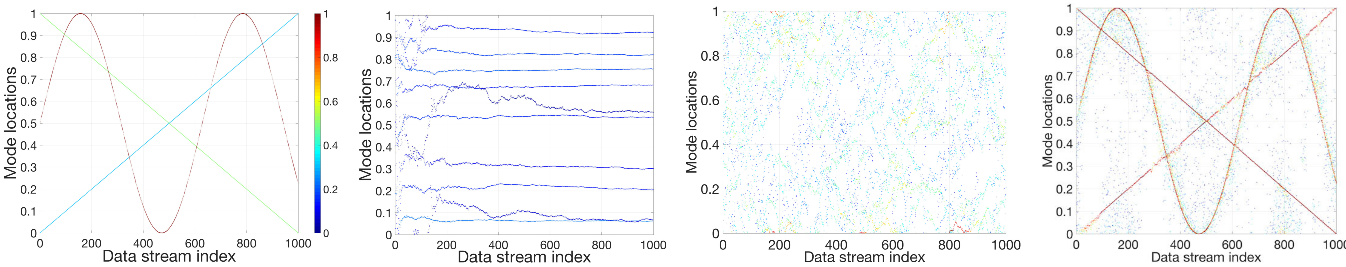

(a) Ground truth (b) GROUSE (c) PETRELS (d) SIST

Fig. 2. Tracking of mode changes in direction-of-arrival estimation under fast yet smooth changes, using GROUSE, PETRELS and SIST: the true mode locations at each stream index is shown in (a), the estimated mode locations at each stream index for 10 modes are shown in (b), (c) and (d) for GROUSE, PETRELS and SIST respectively. All changes are identified and tracked successfully by SIST, but not by GROUSE and PETRELS.

space tracking: The missing data case," *Proceedings of the IEEE*, vol. 106, no. 8, pp. 1293–1310, 2018.

[3] L. Balzano, R. Nowak, and B. Recht, "Online identification and tracking of subspaces from highly incomplete information," in *Proceedings of Allerton Conference on Communication, Control, and Computing*. IEEE, 2010, pp. 704–711.

[4] Y. Chi, Y. C. Eldar, and R. Calderbank, "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.

[5] M. Mardani, G. Mateos, and G. B. Giannakis, "Subspace learning and imputation for streaming big data matrices and tensors," *IEEE Transactions on Signal Processing*, vol. 63, no. 10, pp. 2663–2677, 2015.

[6] H. Guo, C. Qiu, and N. Vaswani, "An online algorithm for separating sparse and low-dimensional signal sequences from their sum," *IEEE Transactions on Signal Processing*, vol. 62, no. 16, pp. 4284–4297, 2014.

[7] Y. Xie, J. Huang, and R. Willett, "Change-point detection for high-dimensional time series with missing data," *IEEE Journal of Selected Topics in Signal Processing*, vol. 7, no. 1, pp. 12–27, 2013.

[8] E. J. Candès and B. Recht, "Exact matrix completion via convex optimization," *Foundations of Computational mathematics*, vol. 9, no. 6, p. 717, 2009.

[9] Y. Li and Y. Chi, "Off-the-grid line spectrum denoising and estimation with multiple measurement vectors," *IEEE Transactions on Signal Processing*, vol. 64, no. 5, pp. 1257–1269, 2016.

[10] Y. Hua and T. K. Sarkar, "Matrix pencil method for estimating parameters of exponentially dampled/undampled sinusoids in noise," *IEEE Transactions on Acoustics, Speech and Signal Processing*, vol. 38, no. 5, pp. 814–824, 1990.

[11] Y. Chen and Y. Chi, "Robust spectral compressed sensing via structured matrix completion," *IEEE Transactions on Information Theory*, vol. 60, no. 10, pp. 6576–6601, 2014.

[12] W. Xu, J. Yi, S. Dasgupta, J.-F. Cai, M. Jacob, and M. Cho, "Separation-free super-resolution from compressed measurements is possible: an orthonormal atomic norm minimization approach," *arXiv preprint arXiv:1711.01396*, 2017.

[13] J.-F. Cai, T. Wang, and K. Wei, "Fast and provable algorithms for spectrally sparse signal reconstruction via low-rank hankel matrix completion," *Applied and Computational Harmonic Analysis*, vol. 46, no. 1, pp. 94 – 121, 2019.

[14] H. Kasai, "Fast online low-rank tensor subspace tracking by CP decomposition using recursive least squares from incomplete observations," *arXiv preprint arXiv:1709.10276*, 2017.

[15] R. Roy and T. Kailath, "Esprit-estimation of signal parameters via rotational invariance techniques," *IEEE Transactions on acoustics, speech, and signal processing*, vol. 37, no. 7, pp. 984–995, 1989.