Streaming Low-Rank Matrix Data Assimilation and Change Identification

Henry Shaowu Yuchi, Matthew Repasky, Terry Ma, Yao Xie
H. Milton Stewart School of Industrial and Systems Engineering
Georgia Institute of Technology
Atlanta, Georgia 30318

2023 57th Asilomar Conference on Signals, Systems, and Computers | 979-8-3503-2574-4/23/$31.00 ©2023 IEEE | DOI: 10.1109/IEEECONF59524.2023.10476834

***Abstract*—Matrix data has been applied to a wide variety of scientific and engineering problems, including image processing, recommendation systems, and network modeling. Much effort has been on modeling individual matrix data using optimization and factorization. However, there has been relatively little research on modeling sequential and streaming observations for matrix data. To this end, we propose a novel Bayesian assimilation approach to model sequential low-rank matrix observations for online anomaly identification. By exploiting the posterior distributional properties of low-rank matrix subspaces, we track the subspace evolution and carry out uncertainty quantification with a computationally efficient procedure. It also enables identifying the changes in streaming matrices in real-time. We evaluate the performance of our method via an application of change identification in piezoresponse force microscopy data for material science studies.**

***Index Terms*—Streaming data, matrix assimilation, uncertainty quantification**

I. Introduction

The inference of matrices from partial observations - *matrix completion* - has garnered significant interest [1]–[4]. It has been applied in numerous real-world problems, ranging from image and video processing [5], recommendation systems [6], and wireless networks [7]. Many matrix completion methods adopt optimization techniques to construct low-rank matrix approximates from incomplete and noisy observations, where matrix nuclear norm minimization is often utilized as a convex relaxation for the low-rank constraint [4], [8], [9]. Other works look into completion via matrix factorization or decomposition [10], including modeling factor matrices in a Bayesian framework [11], [12]. Such Bayesian frameworks enable uncertainty quantification (UQ) on the completed matrix. Furthermore, studying matrix subspaces provides insight into the low-rank matrix modeled. The singular matrix-variate Gaussian (SMG) distribution [13] provides a way to parametrize matrix row and column subspaces. It has been utilized in a Bayesian framework for low-rank matrix completion [14].

Matrix completion generally focuses on estimating the missing entries for a single matrix, so modeling and completing the sequential matrix data calls for dedicated approaches to streaming matrices. Streaming principal component analysis (PCA) is proposed to model the streaming matrix data efficiently in an online setting [15], where a low-dimensional subspace is obtained to approximate higher-dimensional samples. See also [16], [17]. Tracking a factor matrix or matrix subspace in streaming data is proposed by [18], where the principal row and column vectors are updated iteratively to provide completion in streaming matrices. [19] provides a review of related works. However, these works focus on modeling the matrix data stream efficiently for completion or approximation and do not address uncertainty quantification (UQ) in such estimates that are critical for further analysis, e.g., anomaly detection. This is related to high-dimensional sequential change point detection with applications in power systems [20] and sensor networks [21], [22].

To this end, we propose a sequential low-rank matrix assimilation framework that can provide both UQ on matrix subspaces holding feature information and complete the missing matrix entries. To carry out the assimilation, we extend the Bayesian framework of [14] using the SMG distribution to directly model the matrix subspaces in the data stream. Subsequently, we utilize the UQ provided by the assimilation framework to construct a local change identification method via the structure of the low-rank matrix subspace, which is related to [23]–[25].

II. Streaming Matrix Data

Our objective is to model a series of data matrices sequentially, where each is only partially observed and corrupted by random noise. Suppose we have a sequence of observations in matrix data $\{\mathbf{Y}^{(t)}\}_{t=1}^{T} := \{\mathbf{Y}^{(1)}, \mathbf{Y}^{(2)}, \ldots, \mathbf{Y}^{(T)}\}$, where each element of the set $\mathbf{Y}^{(t)} \in \mathbb{R}^{m_1 \times m_2}$ for $t = 1, 2, \ldots, T$, and m_1 and m_2 denote the matrix dimensions. Each matrix $\mathbf{Y}^{(t)}$ is a partial and noisy observation of the underlying low-rank true matrix $\mathbf{X}^{(t)}$ of rank $r \ll \min(m_1, m_2)$. For each matrix, the observed entries form an index set $\Omega_t \subset [m_1] \times [m_2]$ which changes with t, i.e., the observed index set changes between observations. It can then be expressed as follows:

$$Y_{i,j}^{(t)} = X_{i,j}^{(t)} + \epsilon_{i,j}^{(t)} \quad (i,j) \in \Omega^{(t)}. \tag{1}$$

Here, $Y_{i,j}^{(t)}$ is the t-th observation at entry (i, j) corrupted by a Gaussian noise $\epsilon_{i,j}^{(t)}$. We assume $\epsilon_{i,j}^{(t)} \sim \mathcal{N}(0, \eta^2)$, i.e., the noise in each entry follows an i.i.d. zero-mean Gaussian distribution with a fixed variance η^2. In this work, we assume the row and column space of the true matrix $\mathbf{X}^{(t)}$ does not change (or only changes gradually) throughout the observed sequence. This makes streaming data assimilation meaningful, which is often necessary when the dimension of the matrix is large and completion/data storage is no longer efficient so the matrix has to be broken into a sequence. In other words, the low-rank

subspaces (both row and column subspaces) of the matrix can be regarded as fixed until the change point. The subsequent change pattern can be expressed as in Eqn. (2).

$$\begin{aligned} \mathbf{X}^{(\mathbf{t})} &= \mathbf{U}_1\mathbf{D}^{(t)}\mathbf{V}_1^T \qquad t = 1, 2, \ldots, T; \\ \mathbf{X}^{(\mathbf{t})} &= \mathbf{U}_2\mathbf{D}^{(t)}\mathbf{V}_2^T \qquad t > T, \end{aligned} \tag{2}$$

where $\mathbf{U}_1, \mathbf{U}_2 \in \mathbb{R}^{m_1 \times r}$ and $\mathbf{V}_1, \mathbf{V}_2 \in \mathbb{R}^{m_2 \times r}$ are orthogonal matrices of rank r that denote the column and row subspaces of the matrix respectively. ($\mathbf{U}_1^T\mathbf{U}_1 = \mathbf{I}$ and $\mathbf{V}_1^T\mathbf{V}_1 = \mathbf{I}$, and similarly for $\mathbf{U}_2$ and $\mathbf{V}_2$.)

The challenges are two-fold. We first need to track the underlying subspaces of the true matrix $\mathbf{X}^{(\mathbf{t})}$, which is necessary to recover the missing entries and provide UQ. We then need to utilize the UQ derived from the matrix subspaces to identify anomalies in the data stream when the underlying matrix subspace abruptly changes.

III. Sequential Bayesian Subspace Assimilation

A. Model specification

The SMG distribution is defined by [13] as follows: let $\mathbf{Z} \in \mathbb{R}^{m_1 \times m_2}$ be a random matrix with each entry following a Gaussian distribution $Z_{i,j} \sim \mathcal{N}(0, \sigma^2)$ for each $(i, j) \in [m_1] \times [m_2]$. Then, a random matrix $\mathbf{X}$ follows the singular matrix-variate Gaussian distribution if $\mathbf{X} \overset{d}{=} \mathcal{P}_\mathcal{U}\mathbf{Z}\mathcal{P}_\mathcal{V}$ for some projection matrices $\mathcal{P}_\mathcal{U} = \mathbf{U}\mathbf{U}^T$ and $\mathcal{P}_\mathcal{V} = \mathbf{V}\mathbf{V}^T$, with orthogonal matrices $\mathbf{U} \in \mathbb{R}^{m_1 \times r}$, and $\mathbf{V} \in \mathbb{R}^{m_2 \times r}$.

Suppose matrix $\mathbf{X}$ indeed follows the SMG distribution with independent uniform priors $\mathcal{P}_\mathcal{U} \sim U(\mathcal{G}_{r,m_1-r})$, $\mathcal{P}_\mathcal{V} \sim U(\mathcal{G}_{r,m_2-r})$ and fixed σ^2 and r. (Here $\mathcal{G}_{r,m_2-r}$ denotes a *Grassmann manifold*, the space of r-planes in $\mathbb{R}^{m_2}$.) Let $\mathbf{X} = \mathbf{U}\mathbf{D}\mathbf{V}^T$ be the singular value decomposition (SVD) of $\mathbf{X}$, with singular values $\text{diag}(\mathbf{D}) = (d_k)_{k=1}^r$ not necessarily in decreasing order. We find that:

(1) The singular vectors $\mathbf{U}$ and $\mathbf{V}$ follow von Mises-Fisher distributions [26], which are denoted by $\mathcal{MF}(m_1, r, \mathbf{0})$ and $\mathcal{MF}(m_2, r, \mathbf{0})$, respectively.

(2) The singular values $\text{diag}(\mathbf{D}) = (d_k)_{k=1}^r$ follow the repulsed normal distribution, with density:

$$\frac{1}{Z_r(2\pi\sigma^2)^{r/2}} \exp\left\{-\frac{1}{2\sigma^2}\sum_{k=1}^r d_k^2\right\} \prod_{\substack{k,l=1 \\ k<l}}^r |d_k^2 - d_l^2|, \tag{3}$$

where $d_k > 0$ for $k = 1, \ldots, r$. The first finding enables a representation of the matrix by row and column subspaces $\mathbf{U}$ and $\mathbf{V}$ directly, while the second provides the posterior distribution for $\mathbf{D}$ [27]. Using these properties, we obtain the BayeSMG model structure in Table I, where we utilize inverse gamma priors for σ^2 and η^2. It provides us with closed-form posterior conditional distributions for the matrix subspace $\mathbf{U}$ and $\mathbf{V}$. See [14] for more details.

B. Sequential matrix assimilation

To start with, we assign uniform prior distributions for the two subspaces projection matrices in the Grassmann manifold $\mathcal{P}_\mathcal{U} \sim \mathcal{U}[\mathcal{G}_{r,m_1-r}]$ and $\mathcal{P}_\mathcal{V} \sim \mathcal{U}[\mathcal{G}_{r,m_2-r}]$. It is equivalent to the uniform prior distributions for the two subspaces in the Stiefel manifold, $\mathbf{U} \sim \mathcal{U}[\mathcal{V}_{m_1,r}]$ and $\mathbf{V} \sim \mathcal{U}[\mathcal{V}_{m_2,r}]$. The sequential matrix assimilation can hence be conducted iteratively for $t = 1, 2, \ldots, T$. For each observed matrix $\mathbf{Y}^{(t)}$, we carry out the following steps:

TABLE I
Model specification for BayeSMG.

Model	Distribution
Observations	$[\mathbf{Y}_\Omega \vert \mathbf{X}, \eta^2]$: $Y_{i,j} \sim \mathcal{N}(X_{i,j}, \eta^2)$
Low-rank matrix	$[\mathbf{X} \vert \mathbf{U}, \mathbf{V}, \sigma^2] : \mathbf{X} = \mathbf{U}\mathbf{D}\mathbf{V}^T$ $\text{diag}\{\mathbf{D}\} \sim \mathcal{RN}(\mathbf{0}, \sigma^2)$
Priors	
Matrix subspaces	$[\mathcal{P}_\mathcal{U}] \sim \mathcal{U}(\mathcal{G}_{R,m_1-R})$ $[\mathcal{P}_\mathcal{V}] \sim \mathcal{U}(\mathcal{G}_{R,m_2-R})$
Matrix variance	$[\sigma^2] \sim IG(\phi_{\sigma^2}, \beta_{\sigma^2})$
Noise variance	$[\eta^2] \sim IG(\phi_{\eta^2}, \beta_{\eta^2})$

Step 1: subspace estimation. We conduct the singular value decomposition (SVD) on the observed matrix $\mathbf{Y}^{(t)}$ after carrying out a step of low-rank matrix completion (nuclear-norm minimization can be used here, for example).

$$\mathbf{Y}^{(t)} = \hat{\mathbf{U}}^{(t)}\hat{\mathbf{D}}^{(t)}\hat{\mathbf{V}}^{(t)T} + \hat{\mathbf{E}}^{(t)}, \tag{4}$$

where both subspace estimates $\hat{\mathbf{U}}^{(t)}$ and $\hat{\mathbf{V}}^{(t)}$ are constrained by the rank r. The resultant error forms the matrix $\hat{\mathbf{E}}^{(t)}$. We estimate the noise variance $\hat{\eta}^{(t)2}$ from this error matrix directly by calculating the entry-wise deviation.

Step 2: subspace update. Since the matrix subspaces $\mathbf{U}$, $\mathbf{V}$, the noise variance η^2, and the diagonal matrix $\mathbf{D}$ are all estimated in the first step, we can update the posterior conditional distribution for the subspace variable $\mathbf{U}$ which is a von Mises-Fisher distribution.

$$\mathbf{U}|\mathbf{Y}^{(1:t)}, \hat{\mathbf{D}}^{(1:t)}, \hat{\mathbf{V}}^{(1:t)} \propto \exp\left(\text{trace}\left(\sum_{i=1}^t \frac{(\mathbf{Y}^{(i)}\hat{\mathbf{V}}^{(i)}\hat{\mathbf{D}}^{(i)})^T\hat{\mathbf{U}}^{(t)}}{\hat{\eta}^{(i)2}}\right)\right) \tag{5}$$

Similarly, we can derive the updated posterior for $\mathbf{V}$:

$$\mathbf{V}|\mathbf{Y}^{(1:t)}, \hat{\mathbf{D}}^{(1:t)}, \hat{\mathbf{U}}^{(1:t)} \propto \exp\left(\text{trace}\left(\sum_{i=1}^t \frac{(\mathbf{Y}^{(i)T}\hat{\mathbf{U}}^{(i)}\hat{\mathbf{D}}^{(i)})^T\hat{\mathbf{V}}^{(t)}}{\hat{\eta}^{(i)2}}\right)\right) \tag{6}$$

This step enables us to obtain the updated distributions for both subspace variables after seeing a new observation $\mathbf{Y}^{(t)}$. It facilitates the assimilation of the streaming matrix as we only need to store the parameters for the diagonal matrix $\hat{\mathbf{D}}^{(t)}$ and two posterior von Mises-Fisher distributions in Eqn. (5) and (6). They enable tracking the continuous evolution of the matrix subspace features. Furthermore, the uncertainty of the subspaces can be obtained by drawing a few samples from the obtained distributions.

C. Local change identification via UQ

To identify abrupt changes in the matrix subspace structure in the sequence, we compare the estimated subspace variables $\hat{\mathbf{U}}$ and $\hat{\mathbf{V}}$ between the first and second steps in the assimilation procedure. Since we have the assimilated distributions for the two variables from the previous iterations, we can obtain their

respective uncertainty via multiple samples of $\mathbf{U}_j^{(t-1)}$ and $\mathbf{V}_j^{(t-1)}$, where j denotes the index of N realizations drawn for $\mathbf{U}$ and $\mathbf{V}$ which can be efficiently sampled [28]. (In this work N=500.) Comparing the subspaces, we calculate the Frobenius norms $\|\cdot\|_F$, a common distance metric in Grassmann manifolds [29]. The confidence interval (CI) can therefore be approximated using the samples. The CI for $\mathbf{U}^{(t-1)}$ as a necessary condition can be expressed as

$$\Pr\left(0 \leq \|\mathbf{U}^{(t-1)}\mathbf{U}^{(t-1)T} - \hat{\mathbf{U}}^{(t-1)}\hat{\mathbf{U}}^{(t-1)T}\|_F \leq \gamma\right) = (1-\alpha)\times 100\%, \quad \text{where } \gamma \approx Q(1-\alpha) \text{ over } \{\mathbf{U}_j^{(t-1)}\}_{j=1}^N. \tag{7}$$

Here Q denotes the quantile function and $0 < \alpha < 1$ denotes the significance level. Same goes with $\mathbf{V}^{(t-1)}$. The CI is then treated as the change statistics.

$$\begin{aligned} &\frac{1}{N}\sum_{j=1}^{N}\|\mathbf{U}_j^{(t-1)}\mathbf{U}_j^{(t-1)T} - \hat{\mathbf{U}}^{(t)}\hat{\mathbf{U}}^{(t)T}\|_F \leq b_1; \\ &\frac{1}{N}\sum_{j=1}^{N}\|\mathbf{V}_j^{(t-1)}\mathbf{V}_j^{(t-1)T} - \hat{\mathbf{V}}^{(t)}\hat{\mathbf{V}}^{(t)T}\|_F \leq b_2. \end{aligned} \tag{8}$$

Both b_1 and b_2 in Eqn. 8 are thresholds that can be specified. When either threshold is breached, we declare a change has taken place with the matrix subspace structure altered significantly.

IV. Numerical Simulation

We first analyze the efficacy of the proposed matrix assimilation framework by applying it to a numerical simulation experiment. In this simulation example, we look into a sequence of 100 matrices of dimension 64×64, denoted by $\mathbf{Y}^{(t)} \in \mathbb{R}^{64\times 64}$, where $t = 1, 2, \ldots, 100$. Each matrix in the sequence is generated using the following expression:

$$\mathbf{Y}^{(t)} = \mathbf{X}^{(t)} + \mathbf{E}^{(t)} = \mathbf{U}\mathbf{D}^{(t)}\mathbf{V}^T + \mathbf{E}^{(t)}, \tag{9}$$

where $\mathbf{D}^{(t)}$ is a diagonal matrix with each diagonal entry drawn uniformly at random, and $\mathbf{E}^{(t)}$ is the noise matrix. The true matrix at each time step $\mathbf{X}^{(t)}$ is standardized, while each entry of $\mathbf{E}^{(t)}$ is drawn independently from a Gaussian distribution $\mathcal{N}(0, 0.05^2)$. However, the subspace variables $\mathbf{U}$ and $\mathbf{V}$ are fixed from $t = 0$ and only change once at the change point $\tau = 76$. The subspace variables before and after the change point are randomly generated as well. We set the rank of the matrix sequence to $R = 5$; therefore, both subspace variables are of size 100×5, which is of low rank. Finally, we mask the matrix observations by removing the observations of 25% of the matrix entries uniformly at random at each t, precipitating missing entries to the matrix stream.

To evaluate the proposed method, we compare the proposed change identification scheme against two other methods. The first baseline carries out change detection utilizing principal component analysis (PCA), which is commonly used when dealing with matrix data. For every matrix in the sequence, PCA is applied to decompose the matrix after imputing 0's to the missing entries, if any, expressed as $\mathbf{Y}^{(t)} = \mathbf{W}^{(t)}\mathbf{T}^{(t)}$, where $\mathbf{T}^{(t)}$ denotes the principal directions and $\mathbf{W}^{(t)}$ denotes the weights. The computation is done via singular value decomposition. We then truncate $\mathbf{T}^{(t)}$ to obtain the first R principal directions denoted $\hat{\mathbf{T}}^{(t)} \in \mathbb{R}^{m_1\times R}$. The change of estimated $\hat{\mathbf{T}}^{(t)}$ is tracked and treated as the test statistic. The rule can be expressed as

$$\|\hat{\mathbf{T}}^{(t)}\hat{\mathbf{T}}^{(t)T} - \hat{\mathbf{T}}^{(t-1)}\hat{\mathbf{T}}^{(t-1)T}\|_F \leq b_{PCA}. \tag{10}$$

When the principal directions change significantly in an abrupt manner, we report there is a pattern change in the data stream.

The second baseline carries out change identification by looking into the change of matrix entries between steps. When observing the matrix at each time step t, we compare the entries of the current step against those in the last step which are observed in both steps. We then calculate the percentage change of the entry values and subsequently mark the entries that have changed more than 90% between the two steps. We use the proportion of the number of entries marked against the total number of matrix entries as the test statistic. We call this scheme Absolute Entry Difference (AED) change identification with the threshold denoted as b_{AED} in a similar manner to (10).

In this simulation and the experiments that follow, we compare the proposed method against these two baselines by evaluating their accuracy in detecting the actual change in the sequence as in mean absolute error (MAE) between the true change point τ and their estimates $\hat{\tau}$ over 25 replications of the whole procedure, i.e., $|\tau - \hat{\tau}|$.

TABLE II
Results for Numerical Simulation

Error Metric	Assimilation	PCA	AED
MAE	0.28	0.92	0.92

For this simulation, we set the threshold variables by looking at the collection of test statistics as the matrix sequences are analyzed. For b_1, b_2, b_{PCA}, and b_{AED}, we set them as the sum of the running mean of the test statistics (denoted by μ) and three times the running deviation (denoted by δ). In other words, for each method, if the new test statistic exceeds $\mu + 3\delta$ of all previous test statistics already obtained, a change of pattern is declared. The results over 25 replications are obtained as in Table II. We can see the errors in detection are quite small for all three methods, however, the MAE by the proposed method is much smaller than the other two by roughly 70%, which is significant. It can be concluded that the matrix assimilation method is quicker to identify a subspace change in the noisy and partially observed matrix stream.

V. Application: Image Change Detection

Subsequently, we apply the proposed matrix assimilation framework to streaming image data. The intensity values of the pixels of gray-scale images are treated as values of individual matrix entries, so a series of sequential images can be approached as streaming matrices. In this experiment, we utilize our sequential subspace assimilation method to identify underlying changes in images. Additionally, all matrix entries are observed with noise in this particular application.

We look into a series of images taken to track the activity of the sun. The image stream contains 300 screens of pictures taken, which happen to have captured the occurrence of a solar flare incident as shown in Fig. 1. Each image can be converted to a 232×232 matrix, while the solar flare incident takes place at $\tau = 42$ by visual inspection. The incident led to a significant change in the visual observation of the sun. Since human eyes are less susceptible to high-frequency components in images, we need only track the lower frequency features, which leads to the matrix being low rank and our selecting $R = 10$. The images are corrupted with a Gaussian noise $\mathcal{N}(0, 0.05^2)$ after standardization.

We carry out the analysis using the assimilation method for ten replications and obtain the MAE as the error metric, with results shown in Table III. The average errors/detection delays by our proposed method are smaller than the other two baselines by approximately 60%. This suggests that the matrix assimilation method can efficiently tell the change in the image subspace in a consistent manner.

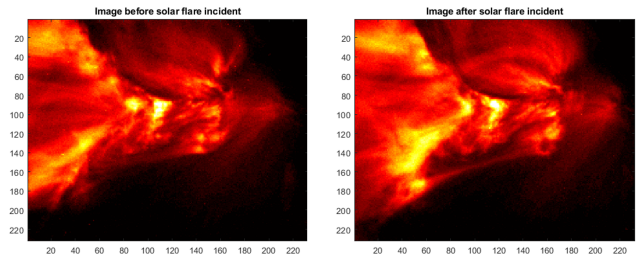

Fig. 1. Illustration of the images before and after the solar flare incident.

TABLE III
RESULTS FOR SOLAR FLARE IMAGE EXAMPLE

Error Metric	Assimilation	PCA	AED
MAE	6.1	15.0	15.0

VI. APPLICATION: POLARIZATION SWITCHING IN PFM

Finally, we apply our proposed matrix assimilation framework to a sequence of piezoresponse force microscopy measurement (PFM) data. PFM is a scanning probe microscopy technique used to study the nanoscale characteristics of ferroelectric materials [30]. It is a powerful tool for high-resolution imaging, manipulation, and spectroscopic measurements [31]. Although PFM is effective in characterizing ferroelectric materials, it is not without challenges. Some of the frequency response signals are of low signal-to-noise ratio (SNR). It is known this leads to poor physical interpretation, as they are deemed unreliable and cannot be used. Moreover, low SNR is more likely to take place during polarization switching, which is of vital interest to study changes in material electromechanical properties.

Using PFM, we obtain the amplitude components of the spectral response within a fixed range of frequencies (250 probing frequency values within a 20 kHz range). A scan is done with a set of different electric field inputs applied to the material surface by PFM as illustrated in Figure 2. The spectral response across different inputs can then be stacked together to form a data matrix. The electric fields change their magnitude and direction gradually between each scanning, giving us a series of matrices. The frequency response is of low SNR when the electrical field is weak, causing unreliable matrix entries that need to be discarded. This gives us matrices with missing entries. Additionally, polarization switching refers to the magnitude of the electric field approaching zero and changing direction. It leads to more missing entries while the data matrix abruptly changes its structure as the material property alters. We aim to both recover the missing entries before polarization switching and identify the switching when it takes place.

From the PFM experiment, 145 data matrices are obtained, each of which has dimension 100×250, reflecting the spectral response from 250 probing frequencies and over 100 different electric field input settings. Polarization switching takes place during the latter half of the experiment, which we aim to detect using our proposed framework. The data matrices before and after the change point are shown in Fig 3. When approaching polarization switching, almost entire rows of the matrix are declared missing, which causes trouble for direct completion for an individual matrix. However, the underlying matrix subspace does not change significantly until the switching, enabling us to utilize the proposed assimilation procedure to estimate matrix subspaces for carrying out completion and change identification.

Again, we compare our proposed strategy to the two baselines. For efficiency, we set $R = 3$ for the assimilation method. We replicate the three methods on the data 20 times and calculate the MAE in detecting the polarization switching. The results are shown in Table IV.

TABLE IV
RESULTS FOR PFM APPLICATION

Error Metric	Assimilation	PCA	AED
MAE	8.7	23.8	28.0

The results clearly show our proposed method performs much better than the other two baselines. Note the SNR is very low when approaching the change point and coincides with many missing entries present in the data stream, so it is expected to be challenging to capture the change swiftly. This led to a slower detection for all three methods compared to the previous application, but the assimilation method still achieved a 60-70% lower detection error.

VII. CONCLUSIONS

We have in this work proposed a novel sequential Bayesian matrix assimilation approach to model sequential low-rank matrix observations and provide uncertainty quantification for the underlying subspace features in the streaming data. Furthermore, it utilizes the UQ to carry out local anomaly identification. Through the numerical experiment and the two scientific applications, we have demonstrated the clear potential and better performance of the proposed sequential matrix assimilation framework in both efficiently modeling streaming matrix data and carrying out robust change identification over the comparison methods. It is empowered by uncertainty

quantification for streaming matrix data without having to conduct computationally intensive procedures such as MCMC.

The limitation of the assimilation method is on the computational expenditure in sampling matrix subspaces, which grows with matrix rank and dimension. For future work, we intend to improve this by investigating efficient matrix subspace sampling.

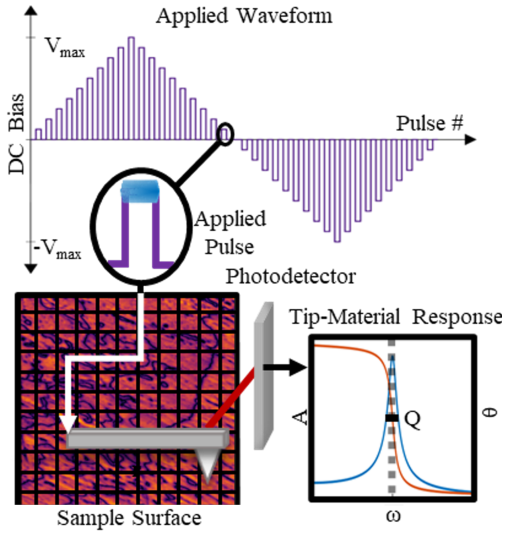

Fig. 2. Illustration of PFM measurements on a grid with the waveform of applied electric field bias.

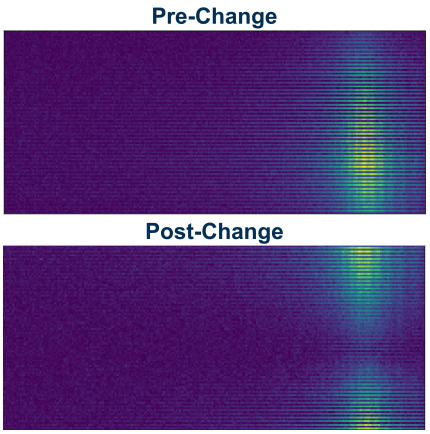

Fig. 3. Visualization of spectral response amplitude components before and after polarization switching.

REFERENCES

[1] E. J. Candès and B. Recht, "Exact matrix completion via convex optimization," *Foundations of Computational mathematics*, vol. 9, no. 6, pp. 717–772, 2009.

[2] E. J. Candès and T. Tao, "The power of convex relaxation: Near-optimal matrix completion," *IEEE Transactions on Information Theory*, vol. 56, no. 5, pp. 2053–2080, 2010.

[3] B. Recht, "A simpler approach to matrix completion.," *Journal of Machine Learning Research*, vol. 12, no. 12, 2011.

[4] V. Koltchinskii, K. Lounici, A. B. Tsybakov, *et al.*, "Nuclear-norm penalization and optimal rates for noisy low-rank matrix completion," *The Annals of Statistics*, vol. 39, no. 5, pp. 2302–2329, 2011.

[5] H. Ji, C. Liu, Z. Shen, and Y. Xu, "Robust video denoising using low rank matrix completion," in *2010 IEEE computer society conference on computer vision and pattern recognition*, pp. 1791–1798, IEEE, 2010.

[6] D. Jannach, P. Resnick, A. Tuzhilin, and M. Zanker, "Recommender systems—beyond matrix completion," *Communications of the ACM*, vol. 59, no. 11, pp. 94–102, 2016.

[7] J. Cheng, Q. Ye, H. Jiang, D. Wang, and C. Wang, "Stcdg: An efficient data gathering algorithm based on matrix completion for wireless sensor networks," *IEEE Transactions on Wireless Communications*, vol. 12, no. 2, pp. 850–861, 2012.

[8] Y. Hu, D. Zhang, J. Ye, X. Li, and X. He, "Fast and accurate matrix completion via truncated nuclear norm regularization," *IEEE transactions on pattern analysis and machine intelligence*, vol. 35, no. 9, pp. 2117–2130, 2012.

[9] S. Gu, Q. Xie, D. Meng, W. Zuo, X. Feng, and L. Zhang, "Weighted nuclear norm minimization and its applications to low level vision," *International journal of computer vision*, vol. 121, no. 2, pp. 183–208, 2017.

[10] Y. Chen, J. Fan, C. Ma, and Y. Yan, "Inference and uncertainty quantification for noisy matrix completion," *Proceedings of the National Academy of Sciences*, vol. 116, no. 46, pp. 22931–22937, 2019.

[11] R. Salakhutdinov and A. Mnih, "Bayesian probabilistic matrix factorization using markov chain monte carlo," in *Proceedings of the 25th international conference on Machine learning*, pp. 880–887, 2008.

[12] P. Alquier *et al.*, "A bayesian approach for noisy matrix completion: Optimal rate under general sampling distribution," *Electronic Journal of Statistics*, vol. 9, no. 1, pp. 823–841, 2015.

[13] A. K. Gupta and D. K. Nagar, *Matrix variate distributions*. Chapman and Hall/CRC, 2018.

[14] H. S. Yuchi, S. Mak, and Y. Xie, "Bayesian uncertainty quantification for low-rank matrix completion," *Bayesian Analysis*, vol. 1, no. 1, pp. 1–28, 2022.

[15] I. Mitliagkas, C. Caramanis, and P. Jain, "Memory limited, streaming pca," *Advances in neural information processing systems*, vol. 26, 2013.

[16] A. C. Gilbert, J. Y. Park, and M. B. Wakin, "Sketched svd: Recovering spectral features from compressive measurements," *arXiv preprint arXiv:1211.0361*, 2012.

[17] C. Boutsidis, D. Garber, Z. Karnin, and E. Liberty, "Online principal components analysis," in *Proceedings of the twenty-sixth annual ACM-SIAM symposium on Discrete algorithms*, pp. 887–901, SIAM, 2014.

[18] S.-Y. Yun, M. Lelarge, and A. Proutiere, "Streaming, memory limited matrix completion with noise," *arXiv preprint arXiv:1504.03156*, 2015.

[19] L. Balzano, Y. Chi, and Y. M. Lu, "Streaming pca and subspace tracking: The missing data case," *Proceedings of the IEEE*, vol. 106, no. 8, pp. 1293–1310, 2018.

[20] Y. C. Chen, T. Banerjee, A. D. Dominguez-Garcia, and V. V. Veeravalli, "Quickest line outage detection and identification," *IEEE Transactions on Power Systems*, vol. 31, no. 1, pp. 749–758, 2015.

[21] Y. Xie and D. Siegmund, "Sequential multi-sensor change-point detection," in *2013 Information Theory and Applications Workshop (ITA)*, pp. 1–20, IEEE, 2013.

[22] V. Raghavan and V. V. Veeravalli, "Quickest change detection of a markov process across a sensor array," *IEEE Transactions on Information Theory*, vol. 56, no. 4, pp. 1961–1981, 2010.

[23] C. Zhang, H. Yan, S. Lee, and J. Shi, "Dynamic multivariate functional data modeling via sparse subspace learning," *Technometrics*, vol. 63, no. 3, pp. 370–383, 2021.

[24] L. Xie, Y. Xie, and G. V. Moustakides, "Sequential subspace change point detection," *Sequential Analysis*, vol. 39, no. 3, pp. 307–335, 2020.

[25] Y. Jiao, Y. Chen, and Y. Gu, "Subspace change-point detection: A new model and solution," *IEEE Journal of Selected Topics in Signal Processing*, vol. 12, no. 6, pp. 1224–1239, 2018.

[26] C. G. Khatri and K. V. Mardia, "The von Mises–Fisher matrix distribution in orientation statistics," *Journal of the Royal Statistical Society: Series B (Methodology)*, vol. 39, no. 1, pp. 95–106, 1977.

[27] J. Shen, "On the singular values of Gaussian random matrices," *Linear Algebra and its Applications*, vol. 326, no. 1-3, pp. 1–14, 2001.

[28] P. D. Hoff, "Simulation of the matrix Bingham–von Mises–Fisher distribution, with applications to multivariate and relational data," *Journal of Computational and Graphical Statistics*, vol. 18, no. 2, pp. 438–456, 2009.

[29] L. Qiu, Y. Zhang, and C.-K. Li, "Unitarily invariant metrics on the grassmann space," *SIAM journal on matrix analysis and applications*, vol. 27, no. 2, pp. 507–531, 2005.

[30] A. Gruverman and S. V. Kalinin, "Piezoresponse force microscopy and recent advances in nanoscale studies of ferroelectrics," *Journal of Materials Science*, vol. 41, no. 1, pp. 107–116, 2006.

[31] S. Jesse, H. N. Lee, and S. V. Kalinin, "Quantitative mapping of switching behavior in piezoresponse force microscopy," *Review of Scientific Instruments*, vol. 77, no. 7, p. 073702, 2006.