

ONLINE TENSOR COMPLETION AND FREE SUBMODULE TRACKING WITH THE T-SVD

Kyle Gilman Laura Balzano

University of Michigan
Department of Electrical and Computer Engineering

ABSTRACT

We propose a new online algorithm, called TOUCAN, for the tensor completion problem of imputing missing entries of a low tubal-rank tensor using the tensor-tensor product (t-product) and tensor singular value decomposition (t-SVD) algebraic framework. We also demonstrate TOUCAN's ability to track changing free submodules from highly incomplete streaming 2-D data. TOUCAN uses principles from incremental gradient descent on the Grassmann manifold to solve the tensor completion problem with linear complexity and constant memory in the number of time samples. We compare our results to state-of-the-art batch tensor completion algorithms and matrix completion algorithms. We show our results on real applications to recover temporal MRI data under limited sampling.

Index Terms— t-SVD, t-product, Grassmannian optimization, online tensor completion, MRI reconstruction

1. INTRODUCTION

Modern data is increasingly high-dimensional and multiway, increasing the storage and computational burden of signal processing algorithms. Many practical applications collect data over multiple modalities. Batch processing of large-scale tensor data quickly becomes computationally intractable, and even storing these tensors is problematic as the memory requirements grow rapidly with the number and size of the tensor modes. Additional challenges include large numbers of missing tensor entries and streaming multiway data that needs to be processed on the fly.

In this paper, we consider sampling and recovery of three-way tensors using the algebraic framework of the t-SVD, where tensors are treated as linear operators over the space of oriented matrices [1, 2, 3]. Using this framework, one obtains an SVD-like factorization referred to as the tensor-SVD (t-SVD) with a defined notion of rank referred to as the tubal-rank. A key property of the t-SVD is the optimality of the truncated t-SVD for data approximation under the Frobenius norm measure [4]. The t-SVD has been well-studied in exact tensor recovery [4], image and video inpainting [5, 6, 7], hyperspectral data [8, 9], and solving tensor robust

principal component analysis (RPCA) problems for video foreground/background separation [10, 11].

Most existing t-SVD based methods are batch methods that require all of the data to be stored in memory at computation time and/or compute multiple SVDs, which is time-consuming and inefficient for large-scale data. Little work has been done to extend online matrix completion methods to the case of tensor data using the t-SVD framework, apart from the works in [12] and [11]. The work in [12] must compute multiple SVDs, and the authors in [11] proposed an online tensor RPCA algorithm for free submodule estimation, but their method cannot predict missing tensor values and does not estimate an orthonormal factorization.

We propose a new algorithm called TOUCAN (Tensor Rank-One Update on the Complex Grassmannian) to recover low-rank tensor data from streaming, highly-incomplete multiway data with incremental gradient descent on the product manifold of low-rank matrices in the Fourier domain using the t-SVD. Our method is online by nature, avoids computing the SVD, and scales linearly in computation with the number of samples. We show our method's ability to track dynamically time-varying low-rank FSMs from streaming two-way data in undersampled MRI data. Additionally, our method can be extended to higher-order tensors.

2. PRELIMINARIES

The following notation and preliminaries are adopted from the work in [1, 10]. We denote a three-way tensor as $\mathcal{A}^{n_1 \times n_2 \times n_3}$. We use $\mathcal{A}(i, :, :)$, $\mathcal{A}(:, i, :)$, $\mathcal{A}(:, :, i)$ to denote the tensor's i -th horizontal, lateral, and frontal slices respectively. Frontal slices are also denoted as $\mathcal{A}^{(i)}$. $\mathcal{A}^*$ denotes the conjugate transpose tensor. The conjugate transpose of a tensor $\mathcal{A} \in \mathbb{C}^{n_1 \times n_2 \times n_3}$ is the tensor $\mathcal{A}^* \in \mathbb{C}^{n_2 \times n_1 \times n_3}$ obtained by conjugate transposing each frontal slice of $\mathcal{A}$ and then reversing the order of transposed slices 2 through n_3 .

We denote the Frobenius norm as $\|\mathcal{A}\|_F = \sqrt{\sum_{ijk} |\mathcal{A}_{ijk}|^2}$. Any lateral slice of size $n_1 \times 1 \times n_3$ is denoted $\vec{\mathcal{A}}$, and any tube along the 3rd dimension of length n_3 is denoted as $\vec{\mathcal{A}}$. The matrix conjugate transpose of a matrix A is denoted A' .

Denote the Discrete Fourier Transform (DFT) matrix for operation on a length- n signal as $F_n \in \mathbb{C}^{n \times n}$ and the DFT of some vector $v \in \mathbb{R}^n$ as $\bar{v} = F_n v \in \mathbb{C}^n$. Note that $F_n/\sqrt{n}$ is unitary, i.e. $F_n' F_n = F_n F_n' = nI_n$ and $F_n^{-1} = F_n'/n$. The

Acknowledgements: Supported by AFOSR FA9550-19-1-0026, ARO W911NF1910027, NSF IIS-1838179 and NSF CCF-1845076. LB supported by the IAS Charles Simonyi Endowment.

DFT is commonly computed in $O(n \log n)$ time by the fast Fourier transform (FFT) as $\bar{v} = \text{fft}(v)$.

We denote $\bar{\mathbf{A}} \in \mathbb{C}^{n_1 \times n_2 \times n_3}$ as the result of computing the DFT along the 3rd dimension, i.e. performing the DFT on the tubes of $\mathbf{A}$. The block-diagonal matrix $\bar{A} \in \mathbb{C}^{n_1 n_3 \times n_2 n_3}$ is the $n_3 n_1 \times n_3 n_2$ matrix with n_3 blocks of size $n_1 \times n_2$ that are the frontal slices of $\bar{\mathbf{A}}$, denoted $\bar{A}^{(i)} \quad \forall i = 1, \dots, n_3$, along its diagonal, $\bar{A} = \text{bdiag}(\bar{\mathbf{A}}) = \text{diag}(\bar{A}^{(1)}, \dots, \bar{A}^{(n_3)})$. We define the block-circulant matrix of the frontal slices of $\mathbf{A}$ as

$$\text{bcirc}(\mathbf{A}) = \begin{bmatrix} A^{(1)} & A^{(n_3)} & \dots & A^{(2)} \\ A^{(2)} & A^{(1)} & \dots & A^{(3)} \\ \vdots & \vdots & \ddots & \vdots \\ A^{(n_3)} & A^{(n_3-1)} & \dots & A^{(1)} \end{bmatrix} \in \mathbb{R}^{n_1 n_3 \times n_2 n_3}$$

From properties of block-circulant matrices, $\text{bcirc}(\mathbf{A})$ can be block-diagonalized by the DFT

$$\bar{A} = (F_{n_3} \otimes I_{n_1}) \cdot \text{bcirc}(\mathbf{A}) \cdot (F_{n_3}^{-1} \otimes I_{n_2}) \quad (1)$$

where $\otimes$ denotes the Kronecker product, and $(F_{n_3}^{-1} \otimes I_{n_2}) / \sqrt{n_3}$ is unitary [10]. For $\mathbf{A} \in \mathbb{R}^{n_1 \times n_2 \times n_3}$ we define the unfold operator that maps $\mathbf{A}$ to a matrix of size $n_1 n_3 \times n_2$ and fold as its inverse operator [1]:

$$\begin{aligned} \text{unfold}(\mathbf{A}) &= [A^{(1)'} \quad A^{(2)'} \quad \dots \quad A^{(n_3)'}]' \\ \text{fold}(\text{unfold}(\mathbf{A})) &= \mathbf{A} \end{aligned}$$

[T-product] [1]: Let $\mathbf{A} \in \mathbb{R}^{n_1 \times n_2 \times n_3}$ and $\mathbf{B} \in \mathbb{R}^{n_2 \times l \times n_3}$. The *t-product* $\mathbf{A} * \mathbf{B}$ is defined to be a tensor of size $n_1 \times l \times n_3$,

$$\mathbf{A} * \mathbf{B} = \text{fold}(\text{bcirc}(\mathbf{A}) \cdot \text{unfold}(\mathbf{B})) \quad (2)$$

By considering three-way tensors to be matrices whose entries are a tubes lying in the third dimension, the t-product can be understood as matrix-matrix multiplication but with circular convolution between the matrix elements. This is equivalent to matrix-matrix multiplication in the Fourier domain, i.e. $\mathbf{C} = \mathbf{A} * \mathbf{B}$ is equivalent to $\bar{\mathbf{C}} = \bar{\mathbf{A}} \bar{\mathbf{B}}$, and can be computed using FFTs as shown in Alg. 1 of [5].

[Identity tensor] [1] The identity tensor $\mathbf{I}_{n \times n \times n_3} \in \mathbb{R}^{n \times n \times n_3}$ is the tensor whose first frontal slice is the $n \times n$ identity matrix, and all other frontal slices are zeros. Property: $\mathbf{A} * \mathbf{I} = \mathbf{I} * \mathbf{A} = \mathbf{A}$

[Orthogonal tensor] [1] A tensor $\mathbf{Q} \in \mathbb{R}^{n \times n \times n_3}$ is orthogonal if it satisfies $\mathbf{Q}^* * \mathbf{Q} = \mathbf{Q} * \mathbf{Q}^* = \mathbf{I}$.

[T-SVD] [1] Let $\mathbf{A} \in \mathbb{R}^{n_1 \times n_2 \times n_3}$. Then it can be factorized as $\mathbf{A} = \mathbf{U} * \mathbf{S} * \mathbf{V}^*$ where $\mathbf{U} \in \mathbb{R}^{n_1 \times n_1 \times n_3}$, $\mathbf{V} \in \mathbb{R}^{n_2 \times n_2 \times n_3}$ are orthogonal, and $\mathbf{S} \in \mathbb{R}^{n_1 \times n_2 \times n_3}$ is a tensor whose frontal slices are diagonal matrices. For any $\mathbf{A} \in \mathbb{R}^{n_1 \times n_2 \times n_3}$, the tensor tubal rank-rank_t($\mathbf{A}$) is defined as the number of nonzero singular tubes of $\mathbf{S}$ from the t-SVD, i.e., $\text{rank}_t(\mathbf{A}) = \#\{i : \mathbf{S}(i, i, :) = \mathbf{0}\}$.

We refer the reader to [10] for a detailed description on the t-SVD and its efficient implementation using FFTs.

[Free Module over the commutative ring] [11] Define $\mathbb{M}_{n_3}^{n_1}$, a free module of dimension n_3 over the commutative ring $\mathbb{R}(\mathbb{G}_{n_3})$, to be the set of all 2-D lateral slices of size $n_1 \times$

$1 \times n_3$ [11]. Since for any element $\vec{\mathbf{X}} \in \mathbb{M}_{n_3}^{n_1}$ and coefficient tube $\vec{t} \in \mathbb{R}^{1 \times 1 \times n_3}$ in the commutative ring $\mathbb{R}(\mathbb{G}_{n_3})$, the lateral slice $\vec{\mathbf{Y}} = \vec{\mathbf{X}} * \vec{t}$ is also an element of $\mathbb{M}_{n_3}^{n_1}$, and $\mathbb{M}_{n_3}^{n_1}$ is closed under tubal-scalar multiplication. One can construct a spanning basis $\{\vec{\mathbf{u}}_1, \vec{\mathbf{u}}_2, \dots, \vec{\mathbf{u}}_r\}$ for this module, and we can uniquely represent any element $\vec{\mathbf{v}} \in \mathbb{M}_{n_3}^{n_1}$ as a $r < n_1$ dimensional t-linear combination of the spanning basis with some tubal coefficients $\vec{w}_i, \vec{\mathbf{v}} = \sum_{i=1}^r \vec{\mathbf{u}}_i * \vec{w}_i = \mathbf{U} * \vec{\mathbf{W}}$. Here, $\mathbf{U} \in \mathbb{R}^{n_1 \times r \times n_3}$ and $\vec{\mathbf{W}} \in \mathbb{R}^{r \times 1 \times n_3}$.

3. PROPOSED METHOD

Given n_2 2-D data samples $\vec{\mathbf{X}}_1, \dots, \vec{\mathbf{X}}_{n_2}$ of size $n_1 \times n_3$, we arrange them as lateral slices to make a 3-D tensor $\mathbf{X}$ of size $n_1 \times n_2 \times n_3$. We consider lateral slices arriving sequentially in time and containing missing entries. We learn the spanning r -dimensional free submodule (FSM) of $\mathbb{M}_{n_3}^{n_1}$ of this multi-way streaming data in an online way similar to the GROUSE algorithm [13]. We allow for the FSM to possibly evolve over time. For full details, see [14].

Let $\mathbf{U} \in \mathbb{R}^{n_1 \times r \times n_3}$ be an orthonormal tensor whose r lateral slices span the the FSM of $\mathbb{M}_{n_3}^{n_1}$. At every time t , we observe an incomplete lateral slice $\vec{\mathbf{v}}_t \in \mathbb{M}_{n_3}^{n_1}$ on the indices $\Omega_t \subset \{1, \dots, n_1\} \times \{1, \dots, n_3\}$. A natural global optimization problem with squared ℓ_2 error loss is

$$\begin{aligned} \hat{\mathbf{U}}, \hat{\mathbf{W}} = \underset{\substack{\mathbf{U} \in \mathbb{R}^{n_1 \times r \times n_3} \\ \mathbf{W} \in \mathbb{R}^{r \times T \times n_3}}}{\text{argmin}} \quad & \frac{1}{2n_3 T} \sum_{t=1}^T \|\mathcal{A}_{\Omega_t}(\vec{\mathbf{v}}_t - \mathbf{U} * \vec{\mathbf{W}}_t)\|_F^2 \\ & \text{s.t. } \mathbf{U}^* * \mathbf{U} = \mathbf{I}_{r \times n_3}. \end{aligned} \quad (3)$$

Above, $\mathcal{A}_{\Omega_t}(\cdot)$ is the linear operator that extracts the observed samples in the set Ω_t . We minimize the cost function for each slice $\vec{\mathbf{v}}_t$ with a stochastic gradient descent procedure. Let $L_t(\mathbf{U}, \vec{\mathbf{W}}_t) = \frac{1}{2} \|\mathcal{A}_{\Omega_t}(\vec{\mathbf{v}}_t - \mathbf{U} * \vec{\mathbf{W}}_t)\|_F^2$. Then the objective in (3) is $\frac{1}{n_3 T} \sum_{t=1}^T L_t(\mathbf{U}, \vec{\mathbf{W}}_t)$. We will take a gradient step for each L_t sequentially. We write the objective function using the definition of the t-product in (2) and the block-diagonalization in (1). We therefore have $L_t(\mathbf{U}, \vec{\mathbf{W}}_t) = \frac{1}{2} \|\mathcal{F}_{\Omega_t}(\bar{\mathbf{v}} - \bar{\mathbf{U}} \bar{\mathbf{W}})\|_2^2$ in the Fourier domain, where $\bar{\mathbf{v}} = \text{unfold}(\vec{\mathbf{v}}_t) \in \mathbb{C}^{n_1 n_3}$ and $\bar{\mathbf{w}} = \text{unfold}(\vec{\mathbf{W}}_t) \in \mathbb{C}^{r n_3}$ for convenient notation. Above, P_{Ω_t} is a subsampled identity matrix of size $|\Omega_t| \times n_1 n_3$ that selects the corresponding indices in Ω_t . Finally, $\mathcal{F}_{\Omega_t} = P_{\Omega_t} \cdot (F_{n_3}^{-1} \otimes I_{n_1}) / \sqrt{n_3} \in \mathbb{C}^{|\Omega_t| \times n_1 n_3}$ is the subsampled inverse Fourier transform. Here we have used the notation from Eq (1) for $\bar{\mathbf{U}}$ which is of size $n_1 n_3 \times r n_3$ and gives us another representation of $\bar{\mathbf{U}}$, with the frontal slices of $\bar{\mathbf{U}}$ on the diagonal, with n_3 blocks of size $n_1 \times r$. This representation requires an additional constraint, that the elements not on these diagonal blocks must be zero. We define the convex set $\mathcal{K}$ to represent these matrices (without any orthogonality constraint).

This is a nonconvex problem from the biconvexity between $\bar{\mathbf{U}}$ and $\bar{\mathbf{w}}$ and the orthonormality constraint. The prob-

lem is also separable in each frontal slice in the Fourier domain, and each is a Grassmannian optimization problem where $\bar{U}^{(i)} \sim \mathcal{G}(n_1, r) \forall i = 1, \dots, n_3$, denoting $\bar{U}^{(i)}$ as a point on the Grassmannian—the set of all subspaces of dimension r in $\mathbb{C}^{n_1}$. The Grassmannian is a compact Riemannian manifold, and its geodesics can be explicitly computed [13].

Our optimization approach uses block coordinate descent, holding the components of $\bar{U}$ fixed while optimizing over the weights $\bar{w}$ and vice versa. With $\bar{U}$ we solve $\bar{w} = \arg\min_{\bar{a} \in \mathbb{C}^{r \times n_3}} \frac{1}{2} \|\mathcal{F}_{\Omega_t}(\bar{v} - \bar{U}\bar{a})\|_2^2$ with conjugate gradient descent (CGD) using FFTs.

Next we perform incremental projected gradient descent on the complex Grassmannian in each slice with $\bar{w}$ fixed. We compute the gradient of L_t (written in the Fourier domain) projected onto the set $\mathcal{K}$, and then take a geodesic step [13].

Given the partial derivatives of L_t with respect to the components of $\bar{U}$ [15]:

$$\frac{\partial L_t}{\partial \bar{U}} = -\mathcal{F}'_{\Omega_t} \mathcal{F}_{\Omega_t}(\bar{v} - \bar{U}\bar{w})\bar{w}' = -\mathcal{F}'_{\Omega_t} \mathcal{F}_{\Omega_t} \bar{r} \bar{w}' ; \quad (4)$$

the projected gradient on the Grassmannian in Fourier space is given by [16] $\nabla L_t = \mathcal{P}_{\mathcal{K}}((I - \bar{U}\bar{U}') \frac{\partial L_t}{\partial \bar{U}})$ where $\mathcal{P}_{\mathcal{K}}(\cdot)$ projects the gradient onto the closest point in the set $\mathcal{K}$ which sets the non-block-diagonal entries of the gradient to zero. The gradient of the objective on the Grassmannian then has the form

$$\nabla L_t = \begin{bmatrix} -\bar{\rho}^{(1)} \bar{w}^{(1)'} & & 0 \\ & \ddots & \\ 0 & & -\bar{\rho}^{(n_3)} \bar{w}^{(n_3)'} \end{bmatrix} \in \mathbb{C}^{n_1 n_3 \times r n_3}$$

Here, $\bar{\rho}^{(i)} = (I - \bar{U}^{(i)} \bar{U}^{(i)'}) \bar{r}_{\Omega_t}^{(i)}$, $\bar{r}_{\Omega_t} = \mathcal{F}'_{\Omega_t} \mathcal{F}_{\Omega_t} \bar{r} = \text{unfold}(\text{fft}(\Delta_{\Omega_t}(\bar{\mathcal{R}}), \cdot, 3))$, $\bar{\mathcal{R}} = \bar{\mathcal{V}} - \bar{\mathcal{U}} * \bar{\mathcal{W}}$ and $\Delta_{\Omega_t}(\cdot)$ imputes zeros on the unobserved tensor entries.

A gradient step along the geodesic with tangent vector $-\nabla L_t$ is given by Equation (2.65) in [16] and is a function of the singular values and vectors of ∇L_t [13]. We can express the SVD of ∇L_t as a product of block-diagonal matrices where each element on the diagonal of ∇L_t is itself a rank-one matrix with nonzero singular value $\bar{\sigma}^{(i)} := 2\|\bar{\rho}^{(i)}\| \|\bar{w}^{(i)'}\|$.

From [16], we now find that for $\eta > 0$, a step of length η in the direction $-\nabla L_t$ is given by $\bar{\mathcal{U}}_{t+1} = \bar{\mathcal{U}}_t + \bar{\mathcal{H}}$ where, for $\bar{\rho}^{(i)} = \bar{U}^{(i)} \bar{w}^{(i)}$,

$$\bar{H}^{(i)} = \begin{cases} \left(\sin(\bar{\sigma}^{(i)} \eta) \frac{\bar{\rho}^{(i)}}{\|\bar{\rho}^{(i)}\|} + (\cos(\bar{\sigma}^{(i)} \eta) - 1) \frac{\bar{\rho}^{(i)}}{\|\bar{\rho}^{(i)}\|} \right) \frac{\bar{w}^{(i)'}}{\|\bar{w}^{(i)}\|}, & i = 1, \dots, \lceil \frac{n_3+1}{2} \rceil \\ \text{conj}(\bar{H}^{(i)}), & i = \lceil \frac{n_3+1}{2} \rceil + 1, \dots, n_3 \end{cases} \quad (5)$$

Following from [17], we use a greedy step size η above for each slice $\bar{H}^{(i)}$, $\eta = \arctan(\|\bar{\rho}^{(i)}\|/\|\bar{w}^{(i)}\|)$. Using conjugate symmetry of the Fourier transform, we can save significant time by only computing the matrix-vector multiplications on half of the frontal slices in the Fourier domain [10].

The preceding updates give an efficient algorithm for computing each variable in the Fourier domain with simple, efficient linear algebra operations and fast Fourier transforms. TOUCAN is numerically stable by maintaining orthonormality on the tensor Grassmannian and is constant in memory use, scaling linearly with the number of observed data samples instead of in polynomial-time like batch t-SVD methods. TOUCAN is summarized in Algorithm 1. For the problem of *missing tubes* of data, the optimization problem is entirely block-diagonal and, as the work in [18] showed, is separable in each frontal slice in the Fourier domain. The algorithm is similar to, but more efficient than, Alg. 1; $\bar{\mathcal{W}}_t$ can be solved exactly in closed form using psuedo-inverses in the Fourier domain, and $\bar{\rho}^{(i)}$ in Alg. 1 becomes $\bar{r}^{(i)}$.


Algorithm 1 Tensor rank-One Update on the Complex grassmannian (TOUCAN): Arbitrary Missing Tensor Entries

Require: A series of lateral slices $\bar{\mathcal{V}}_t \in \mathbb{R}^{n_1 \times 1 \times n_3}$, $\forall i = 1, \dots, T$ observed on the indices in Ω_t ; rank parameter r , randomly-initialized $\bar{\mathcal{U}}_0 \in \mathbb{C}^{n_1 \times r \times n_3}$.


- 1: **for** $t = 1$ to T **do**
- 2: Compute $\bar{\mathcal{V}}_t = \text{ifft}(\mathcal{V}_t, \cdot, 3)$
- 3: Estimate weights $\bar{W}_t$ with CGD.
- 4: Compute: $\bar{P} = \bar{U}_t \bar{W}_t$, $\bar{\mathcal{P}} = \text{ifft}(\bar{\mathcal{P}}, \cdot, 3)$
- 5: Compute residual: $\bar{\mathcal{R}} = \bar{\mathcal{V}}_t - \bar{\mathcal{P}}$
- 6: Update subspace: $\bar{\mathcal{U}}_{t+1} = \bar{\mathcal{U}}_t + \bar{\mathcal{H}}$ by Eq. (5)
- 7: Transform: $\bar{\mathcal{U}}_{t+1} = \text{fft}(\bar{\mathcal{U}}_{t+1}, \cdot, 3)$
- 8: Transform: $\bar{\mathcal{W}}_t = \text{ifft}(\bar{\mathcal{W}}_t, \cdot, 3)$
- 9: **end for**
- 10: **return** $\bar{\mathcal{U}}_{t+1}, \bar{\mathcal{W}}_t$, $\forall t = 1, \dots, T$

4. EXPERIMENTAL RESULTS

[Numerical results] We first verify the validity and efficiency of TOUCAN in recovering large-scale missing tensor data synthetically generated from isotropic Gaussian distributions. We compute the t-product of two low-tubal rank tensors to yield a tensor of sizes $n_1 = 200, n_2 = 500, n_3 = 20$ and tubal-rank $r = 3$. We sample 50% of tensor entries/tubes randomly according to a Bernoulli distribution. TOUCAN observes one lateral slice at each time instance, solves the inner CGD step to within a set tolerance ($1e - 9$), and is allowed to process over the entire batch more than once until the desired termination tolerance. We compare against batch tensor completion algorithms in [4, 6] (with improved computational efficiency using conjugate symmetry) and to standard matrix PCA algorithms by matricizing the tensor and computing batch matrix completion [19] and GROUSE [13] on each column of the matricized tensor. We plot the normalized root-mean-squared error (NRMSE) of the recovered tensor to the true tensor by elapsed wall clock time in seconds in Fig. 1(a), terminating each algorithm if its NRMSE is less than $1e - 9$. TOUCAN is competitive with state-of-the-art batch method Tensor Factorization (TCTF) [6] in the case of random entry sampling, and TOUCAN's efficient tubal-sampling



(a) Tensor completion of synthetic t-product generated data. Solid lines are experiments with uniformly random samples, and dashed lines are experiments with uniformly sampled tubes.



(b) Tracking dynamic FSMs of different tubal ranks under random sampling. Here, $n_1 = 50$ and $n_3 = 10$.

Fig. 1: Synthetic numerical experiments.

implementation vastly outperforms each other algorithms.

[Dynamic FSM Tracking] We demonstrate TOUCAN’s ability to track a dynamically changing FSM from streaming multiway data with missing entries. We generate a random orthonormal basis $\mathbf{U}$ for various tubal ranks from an i.i.d. Gaussian distribution and draw 2-D lateral slices by t-product with i.i.d. Gaussian weights. We sample 70% of the entries/tubes at random. At each iteration, for TOUCAN’s estimate $\hat{\mathbf{U}}$, we measure $\|\hat{\mathbf{U}} * \hat{\mathbf{U}}^* - \mathbf{U} * \mathbf{U}^*\|_F / \|\mathbf{U} * \mathbf{U}^*\|_F$. We simulate abrupt FSM changes by randomly reinitializing the underlying FSM every 500 slices. The simulation results in Fig. 1(b) show TOUCAN’s ability to reliably re-estimate each new FSM.

[Streaming dynamic MRI reconstruction] Magnetic resonance imaging (MRI) collects a high-dimensional tensor that is often undersampled due to computational limitations exacerbated by large volumetric and dynamic acquisitions. One successful solution to image reconstruction from limited sampling is low-rank tensor completion [20, 21]. A t-SVD factorization of the spatial frequency-by-time (or k - t space) tensor shows low-tubal-rank structure in the real and complex components [20], and t-SVD algorithms have been shown to be proficient at completing the k - t space tensor for image reconstruction. MRI data can also contain significant motion content and time-varying dynamics such as breathing motion. We employ TOUCAN’s ability to track streaming time-dynamic multiway data to recover the k - t space tensor.

We test TOUCAN against the batch t-SVD algorithms and an online tensor completion algorithm in [22] that estimates a low-dimensional Grassmannian in each mode of a tensor unfolding model, called Sequential Tensor Completion algorithm (STC), to recover undersampled MRI data. We use the invivo myocardial perfusion dataset data from [23] whose tensor is in $\mathbb{C}^{k_x \times k_t \times k_y}$ where $k_x = 190$, $k_y = 90$ and k_t , the number of time samples, is 70. We allow the online methods




Fig. 2: Reconstructed images of invivo myocardial perfusion dataset from 40% of k - t space samples.

	NRMSE		SSIM		Comp. Time (s)	
	Random	Tube	Random	Tube	Random	Tube
Original	0	0	1	1	–	–
Zero-filled	0.5755	0.5354	0.4735	0.6364	–	–
TOUCAN	0.2148	0.2108	0.8637	0.8878	3.193	0.8425
TNN-ADMM	0.1144	0.1132	0.9518	0.9580	17.081	16.493
TCTF	0.2096	0.2018	0.8625	0.8941	15.837	14.240
STC	0.4300	–	0.6593	–	15.388	–

Table 1: Invivo myocardial perfusion experiment statistics.

one pass over the data in a streaming way; TOUCAN learns a tubal-rank 5 FSM, and we set the ranks $r_1 = r_2 = 90$, $r_3 = 5$ for STC. TCTF learns a tubal-rank 5 factorization. STC cannot handle tube-sampled data, so we only test it in the case where arbitrarily random entries are missing. We compute the NRMSE and mean structural similarity index measures (SSIM) of the reconstructed images (shown in Fig. 2, along with the total computation times for each algorithm, which are shown in Table 1). We show TOUCAN can achieve competitive reconstruction measures against the batch algorithms in far-less memory and computation time, and completely outperforms STC.

5. CONCLUSION

In this paper we presented a novel algorithm for online low-tubal-rank tensor completion under the t-SVD algebraic framework. Our method avoids computing SVDs, and only needs to update and store a smaller orthonormal tensor and the lateral slice of weights per iteration, leading to a powerful and efficient online algorithm that scales linearly in computation with the number of samples.

A robust version of TOUCAN for tensor RPCA is of interest for problems like multispectral foreground/background separation. We plan to investigate choices for the number of CGD iterations to balance accuracy with running time. Lastly, our method merits comparison to the work in [12] and other tensor algebraic frameworks like CP and Tucker.

6. REFERENCES

- [1] Misha E. Kilmer and Carla D. Martin, "Factorization strategies for third-order tensors," *Linear Algebra and Its Applications - LINEAR ALGEBRA APPL*, vol. 435, no. 8, pp. 148–172, 2013.
- [2] M. Kilmer, K. Braman, N. Hao, and R. Hoover, "Third-order tensors as operators on matrices: A theoretical and computational framework with applications in imaging," *SIAM Journal on Matrix Analysis and Applications*, vol. 34, no. 1, pp. 148–172, 2013.
- [3] C. Martin, R. Shafer, and B. LaRue, "An order-p tensor factorization with applications in imaging," *SIAM Journal on Scientific Computing*, vol. 35, no. 1, pp. A474–A490, 2013.
- [4] Z. Zhang and S. Aeron, "Exact tensor completion using t-SVD," *IEEE Transactions on Signal Processing*, vol. 65, no. 6, pp. 1511–1526, March 2017.
- [5] Z. Zhang, G. Ely, S. Aeron, N. Hao, and M. Kilmer, "Novel methods for multilinear data completion and denoising based on tensor-SVD," in *2014 IEEE Conference on Computer Vision and Pattern Recognition*, June 2014, pp. 3842–3849.
- [6] P. Zhou, C. Lu, Z. Lin, and C. Zhang, "Tensor factorization for low-rank tensor completion," *IEEE Transactions on Image Processing*, vol. 27, no. 3, pp. 1152–1163, March 2018.
- [7] J. Liu, P. Musialski, P. Wonka, and J. Ye, "Tensor completion for estimating missing values in visual data," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 35, no. 1, pp. 208–220, Jan 2013.
- [8] H. Fan, Y. Chen, Y. Guo, H. Zhang, and G. Kuang, "Hyperspectral image restoration using low-rank tensor recovery," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 10, no. 10, pp. 4589–4604, Oct 2017.
- [9] H. Fan, C. Li, Y. Guo, G. Kuang, and J. Ma, "Spatial-spectral total variation regularized low-rank tensor decomposition for hyperspectral image denoising," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 56, no. 10, pp. 6196–6213, Oct 2018.
- [10] Lu Canyi, Jiashi Feng, Yudong Chen, Wei Liu, Zhouchen Lin, and Shuicheng Yan, "Tensor robust principal component analysis with a new tensor nuclear norm," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. PP, 04 2018.
- [11] Z. Zhang, D. Liu, S. Aeron, and A. Vetro, "An online tensor robust PCA algorithm for sequential 2D data," in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, March 2016, pp. 2434–2438.
- [12] John Pothier, Josh Girson, and Shuchin Aeron, "An algorithm for online tensor prediction," 2015.
- [13] Laura Balzano, Robert D. Nowak, and Benjamin Recht, "Online identification and tracking of subspaces from highly incomplete information," *2010 Allerton*, pp. 704–711, 2010.
- [14] Kyle Gilman and Laura Balzano, "Grassmannian optimization for online tensor completion and tracking in the t-SVD algebra," 2020.
- [15] Kaare Brandt Petersen and Michael Syskind Pedersen, "The matrix cookbook," 2012.
- [16] A. Edelman, T. Arias, and S. Smith, "The geometry of algorithms with orthogonality constraints," *SIAM Journal on Matrix Analysis and Applications*, vol. 20, no. 2, pp. 303–353, 1998.
- [17] Dejiao Zhang and Laura Balzano, "Global convergence of a grassmannian gradient descent algorithm for subspace estimation," in *AISTATS*, 2015.
- [18] J. Girson and S. Aeron, "Tensor completion via optimization on the product of matrix manifolds," in *2015 IEEE 6th International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP)*, Dec 2015, pp. 177–180.
- [19] Emmanuel J. Candès and Benjamin Recht, "Exact matrix completion via convex optimization," *Foundations of Computational Mathematics*, vol. 9, no. 6, pp. 717, Apr 2009.
- [20] D. Banco, S. Aeron, and W. S. Hoge, "Sampling and recovery of MRI data using low rank tensor models," in *2016 38th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, Aug 2016, pp. 448–452.
- [21] 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, May 2015.
- [22] K. Xie, L. Wang, X. Wang, G. Xie, J. Wen, G. Zhang, J. Cao, and D. Zhang, "Accurate recovery of internet traffic data: A sequential tensor completion approach," *IEEE/ACM Transactions on Networking*, vol. 26, no. 2, pp. 793–806, April 2018.
- [23] S. G. Lingala, Y. Hu, E. DiBella, and M. Jacob, "Accelerated dynamic mri exploiting sparsity and low-rank structure: k-t SLR," *IEEE Transactions on Medical Imaging*, vol. 30, no. 5, pp. 1042–1054, May 2011.