A dual framework for low-rank tensor completion

Madhav Nimishakavi[‡]
madhav@iisc.ac.in

Pratik Jawanpuria[†]
pratik.jawanpuria@microsoft.com

Bamdev Mishra[†]
bamdevm@microsoft.com

[‡]Indian Institute of Science
[†]Microsoft

Abstract

One of the popular approaches for low-rank tensor completion is to use the *latent trace norm* regularization. However, most existing works in this direction learn a sparse combination of tensors. In this work, we fill this gap by proposing a variant of the latent trace norm that helps in learning a non-sparse combination of tensors. We develop a dual framework for solving the low-rank tensor completion problem. We first show a novel characterization of the dual solution space with an interesting factorization of the optimal solution. Overall, the optimal solution is shown to lie on a Cartesian product of Riemannian manifolds. Furthermore, we exploit the versatile Riemannian optimization framework for proposing computationally efficient trust region algorithm. The experiments illustrate the efficacy of the proposed algorithm on several real-world datasets across applications.

1   Introduction

Tensors are multidimensional or K-way arrays, which provide a natural way to represent multi-modal data [CPZ+17a, CPZ+17b]. Low-rank tensor completion problem, in particular, aims to recover a low-rank tensor from partially observed tensor [ADKM11]. This problem has numerous applications in image/video inpainting [LMWY13, KSV14], link-prediction [EAC15], and recommendation systems [SNM08], to name a few.

In this work, we focus on trace norm regularized low-rank tensor completion problem of the form

$$\min_{\boldsymbol{\mathcal{W}}\in\mathbb{R}^{n_1\times n_2\times\ldots\times n_K}} \|\boldsymbol{\mathcal{W}}_\Omega - \boldsymbol{\mathcal{Y}}_\Omega\|_F^2 + \frac{1}{\lambda}R(\boldsymbol{\mathcal{W}}), \tag{1}$$

where $\boldsymbol{\mathcal{Y}}_\Omega \in \mathbb{R}^{n_1\times\ldots\times n_K}$ is a partially observed K- mode tensor, whose entries are only known for a subset of indices Ω. $(\boldsymbol{\mathcal{W}}_\Omega)_{(i_1,\ldots,i_K)} = \boldsymbol{\mathcal{W}}_{(i_1,\ldots,i_K)}$, if $(i_1,\ldots,i_K) \in \Omega$ and 0 otherwise, $\|\cdot\|_F$ is the Frobenius norm , $R(\cdot)$ is a low-rank promoting regularizer, and $\lambda > 0$ is the regularization parameter.

Similar to the matrix completion problem, the trace norm regularization has been used to enforce the low-rank constraint for the tensor completion problem. The works [THK10, TS13] discuss the *overlapped* and *latent* trace norm regularizations for tensors. In particular, [TS13, WST14] show that the latent trace norm has certain better tensor reconstruction bounds. The latent trace norm regularization learns the tensor as a *sparse* combination of different tensors. In our work, we empirically motivate the need for learning non-sparse combination of tensors and propose a variant of the latent trace norm that learns a *non-sparse* combination of tensors. We show a novel

characterization of the solution space that allows for a compact storage of the tensor, thereby allowing to develop scalable optimization formulations. Concretely, we make the following contributions in this paper.

- We propose a novel trace norm regularizer for low-rank tensor completion problem, which learns a tensor as a non-sparse combination of tensors. In contrast, the more popular latent trace norm regularizer [THK10, TS13, WST14] learns a highly sparse combination of tensors. Non-sparse combination helps in capturing information along all the modes.

- We propose a dual framework for analyzing the problem formulation. This provides interesting insights into the solution space of the tensor completion problem, e.g., how the solutions along different modes are related, allowing a compact representation of the tensor.

- Exploiting the characterization of the solution space, we develop a fixed-rank formulation. Our optimization problem is on Riemannian spectrahedron manifolds and we propose computationally efficient trust-region algorithm for our formulation.

Numerical comparisons on real-world datasets for different applications such as video and hyperspectral-image completion, link prediction, and movie recommendation show that the proposed algorithm outperforms state-of-the-art latent trace norm regularized algorithms.

The organization of the paper is as follows. Related works are discussed in Section 2. In Section 3, we study a particular trace norm based tensor completion formulation. Theorem 1 shows the characterization of the solution space. Building on this, we show two particular fixed-rank formulations. Both of these problems have the structure of optimization on Riemannian manifolds. The optimization related ingredients and their computational cost are subsequently discussed in Section 4. In Section 5, numerical comparisons on real data sets for three different applications: video and hyperspectral-image completion, link prediction, and movie recommendation, show that our proposed algorithms outperform various state-of-the-art latent trace norm regularized algorithms.

The Matlab codes are available at `https://pratikjawanpuria.com/`.

2   Related work

Trace norm regularized tensor completion formulations. The works [LMWY13, TS13, SDLS14, RpABbP13, CYZ+16] discuss the *overlapped trace norm* regularization for tensor learning. The overlapped trace norm is motivated as a convex proxy for minimizing the *Tucker* (multilinear) rank of a tensor. The overlapped trace norm is defined as: $R(\mathcal{W}) \coloneqq \sum_{k=1}^{K} \|\mathcal{W}_k\|_*$, where $\mathcal{W}_k$ is the mode-k matrix unfolding of the tensor $\mathcal{W}$ [KB09] and $\|\cdot\|_*$ denotes the trace norm regularizer.

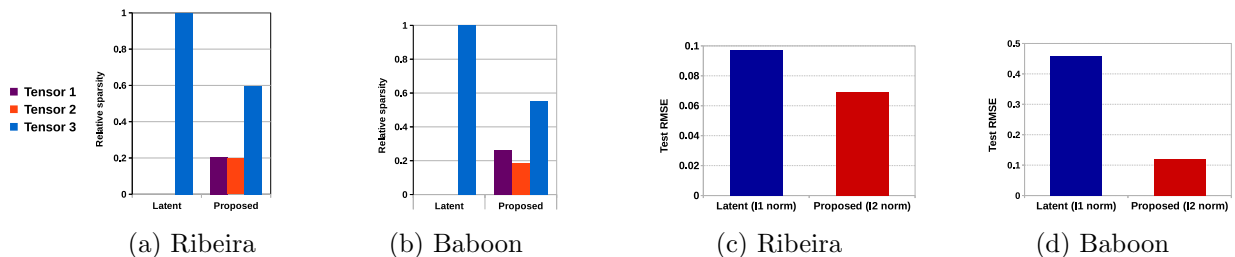| (a) Ribeira | (b) Baboon | (c) Ribeira | (d) Baboon |

Figure 1: (a) & (b) Relative sparsity of each tensor in the mixture of tensors for Ribeira and Baboon datasets. Our proposed formulation learns a ℓ_2-norm based non-sparse combination of tensors; (c) & (d) show that the proposed non-sparse combination obtain better generalization performance on both the datasets.

$\boldsymbol{\mathcal{W}}_k$ is a $n_k \times \Pi_{j \neq k} n_j$ matrix obtained by concatenating mode-k fibers (column vectors) of the form $\boldsymbol{\mathcal{W}}_{(i_1,\ldots,i_{k-1},:,i_{k+1},\ldots,i_K)}$ [KB09].

Latent trace norm is another convex regularizer used for low-rank tensor learning [THK10, TSHK11, TS13, WST14, GYK17]. In this setting, the tensor $\boldsymbol{\mathcal{W}}$ is modeled as *sum* of K (unknown) tensors $\boldsymbol{\mathcal{W}}^{(1)},\ldots,\boldsymbol{\mathcal{W}}^{(K)}$ such that $\boldsymbol{\mathcal{W}}_k^{(k)}$ are low-rank matrices. The latent trace norm is defined as:

$$R(\boldsymbol{\mathcal{W}}) \coloneqq \inf_{\sum_{k=1}^K \boldsymbol{\mathcal{W}}^{(k)} = \boldsymbol{\mathcal{W}};\; \boldsymbol{\mathcal{W}}^{(k)} \in \mathbb{R}^{n_1 \times \ldots \times n_K}} \sum_{k=1}^K \|\boldsymbol{\mathcal{W}}_k^{(k)}\|_*, \tag{2}$$

A variant of the latent trace norm ($\|\boldsymbol{\mathcal{W}}_k^{(k)}\|_*$ scaled by $1/\sqrt{n_k}$) is analyzed in [WST14]. Latent trace norm and its scaled variant achieve better recovery bounds than overlapped trace norm [TS13, WST14]. Recently, [GYK17] proposed a scalable latent trace norm based Frank-Wolfe algorithm for tensor completion.

The latent trace norm (2) corresponds to the sparsity inducing ℓ_1-norm penalization across $\|\boldsymbol{\mathcal{W}}_k^{(k)}\|_*$. Hence, it learns $\boldsymbol{\mathcal{W}}$ as a sparse combination of $\boldsymbol{\mathcal{W}}^{(k)}$. In case of high sparsity, it may result in selecting only one of the tensors $\boldsymbol{\mathcal{W}}^{(k)}$ as $\boldsymbol{\mathcal{W}}$, i.e., $\boldsymbol{\mathcal{W}} = \boldsymbol{\mathcal{W}}^{(k)}$ for some k, in which case $\boldsymbol{\mathcal{W}}$ is essentially learned as a low-rank matrix. In several real-world applications, tensor data cannot be mapped to a low-rank matrix structure and they require a higher order structure. Therefore, we propose a regularizer which learns a non-sparse combination of $\boldsymbol{\mathcal{W}}^{(k)}$. Non-sparse norms have led to better generalization performance in other machine learning settings [CMR09, Suz11, JVN14].

We show the benefit of learning a non-sparse mixture of tensors as against a sparse mixture on two datasets: Ribeira and Baboon (refer Section 5 for details). Figures 1(a) and 1(b) show the relative sparsity of the optimally learned tensors in the mixture as learned by the ℓ_1-regularized latent trace norm based model (2) [TS13, WST14, GYK17] versus the proposed ℓ_2-regularized model (discussed in Section 3). The relative sparsity for each $\boldsymbol{\mathcal{W}}^{(k)}$ in the mixture is computed as $\|\boldsymbol{\mathcal{W}}^{(k)}\|_F / \sum_k \|\boldsymbol{\mathcal{W}}^{(k)}\|_F$. In both the datasets, our model learns a non-sparse combination of tensors, whereas the latent trace norm based model learns a highly skewed mixture of tensors. The proposed non-sparse tensor combination also leads to better generalization performance, as can be observed in the Figures 1(c) and 1(d). In the particular case of Baboon dataset, the latent trace norm essentially learns $\boldsymbol{\mathcal{W}}$ as a low-rank matrix ($\boldsymbol{\mathcal{W}} = \boldsymbol{\mathcal{W}}^{(3)}$) and consequently obtains poor generalization.

Other tensor completion formulations. Other approaches for low-rank tensor completion include tensor decomposition methods like Tucker and CP [KB09, CPZ$^+$17a, CPZ$^+$17b]. They generalize the notion of singular value decomposition of matrices to tensors. Recently, [KSV14] exploits the Riemannian geometry of fixed multilinear rank to learn factor matrices and the core tensor. They propose a computationally efficient non-linear conjugate gradient method for optimization over manifolds of tensors of fixed multilinear rank. [KM16] further propose an efficient preconditioner for low-rank tensor learning with the Tucker decomposition. [ZZC15] propose a Bayesian probabilistic CP model for performing tensor completion. Tensor completion algorithms based on tensor *tubal-rank* have been recently proposed in [ZEA$^+$14, LAAW16].

3 Non-sparse latent trace norm, duality, and novel formulations for low-rank tensor completion

We propose the following formulation for learning the low-rank tensor $\boldsymbol{\mathcal{W}}$

$$\min_{\boldsymbol{\mathcal{W}}^{(k)} \in \mathbb{R}^{n_1 \times \ldots \times n_K}} \|\boldsymbol{\mathcal{W}}_\Omega - \boldsymbol{\mathcal{Y}}_\Omega\|_F^2 + \sum_k \frac{1}{\lambda_k} \left\|\boldsymbol{\mathcal{W}}_k^{(k)}\right\|_*^2, \tag{3}$$

3

where $\boldsymbol{\mathcal{W}} = \sum_k \boldsymbol{\mathcal{W}}^{(k)}$ is the learned tensor. It should be noted that the proposed regularizer in (3) employs the ℓ_2-norm over $\|\boldsymbol{\mathcal{W}}_k^{(k)}\|_*$. In contrast, the latent trace norm regularizer (2) has the ℓ_1-norm over $\|\boldsymbol{\mathcal{W}}_k^{(k)}\|_*$.

While the existing tensor completion approaches [KM16, GYK17, KSV14, LMWY13, TS13, SDLS14] mostly discuss a primal formulation similar to (1), we propose a novel dual framework for our analysis. The use of dual framework, e.g., in the matrix completion problem [XJ12, PTJY10, JM18], often leads to novel insights into the solution space of the primal problem.

We begin by discussing how to obtain the dual formulation of (3). Later, we explain how the insights from the dual framework motivate us to propose *two novel* fixed-rank formulations for (3). As a first step, we exploit the following *variational characterization* of the trace norm studied in [AEP06, Theorem 4.1]. Given $\mathbf{X} \in \mathbb{R}^{d \times T}$, the following result holds:

$$\|\mathbf{X}\|_*^2 = \min_{\Theta \in \mathcal{P}^d, \mathrm{range}(\mathbf{X}) \subseteq \mathrm{range}(\Theta)} \langle \Theta^\dagger, \mathbf{X}\mathbf{X}^\top \rangle, \tag{4}$$

where $\mathcal{P}^d$ denotes the set of $d \times d$ positive semi-definite matrices with *unit trace*, $\Theta^\dagger$ denotes the pseudo-inverse of Θ, $\mathrm{range}(\Theta) = \{\Theta z : z \in \mathbb{R}^d\}$, and $\langle \cdot, \cdot \rangle$ is the inner product. The expression for optimal Θ^* is $\Theta^* = \sqrt{\mathbf{X}\mathbf{X}^\top}/\mathrm{trace}(\sqrt{\mathbf{X}\mathbf{X}^\top})$ [AEP06], and hence the ranks of Θ and $\mathbf{X}$ are equal at optimality. Thus, (4) implicitly transfers the low-rank constraint on $\mathbf{X}$ (due to trace norm) to an auxiliary variable $\Theta \in \mathcal{P}^d$. It is well known that positive semi-definite matrix Θ with unit trace constraint implies the ℓ_1-norm constraint on the eigenvalues of Θ, leading to low-rankness of Θ.

Using the result (4) in (3) leads to K auxiliary matrices, one $\Theta_k \in \mathcal{P}^{n_k}$ corresponding to every $\boldsymbol{\mathcal{W}}_k^{(k)}$ (mode-k matrix unfolding of the tensor $\boldsymbol{\mathcal{W}}^{(k)}$). It should also be noted that $\Theta_k \in \mathcal{P}^{n_k}$ are low-rank matrices. We now present the following theorem that states an equivalent minimax formulation of (3).

Theorem 1. *An equivalent minimax formulation of the problem (3) is*

$$\min_{\Theta_1 \in \mathcal{P}^{n_1}, \dots, \Theta_K \in \mathcal{P}^{n_K}} \max_{\boldsymbol{\mathcal{Z}} \in \mathcal{C}} \ \langle \boldsymbol{\mathcal{Z}}, \boldsymbol{\mathcal{Y}}_\Omega \rangle - \frac{1}{4} \|\boldsymbol{\mathcal{Z}}\|_F^2 - \sum_k \frac{\lambda_k}{2} \langle \Theta_k, \boldsymbol{\mathcal{Z}}_k \boldsymbol{\mathcal{Z}}_k^\top \rangle, \tag{5}$$

where $\boldsymbol{\mathcal{Z}}$ is the dual tensor variable corresponding to the primal problem (3) and $\boldsymbol{\mathcal{Z}}_k$ is the mode-k unfolding of $\boldsymbol{\mathcal{Z}}$. The set $\mathcal{C} := \{\boldsymbol{\mathcal{Z}} \in \mathbb{R}^{n_1 \times \dots \times n_K} : \boldsymbol{\mathcal{Z}}_{(i_1, \dots, i_K)} = 0, (i_1, \dots, i_K) \notin \Omega\}$ constrains $\boldsymbol{\mathcal{Z}}$ to be a sparse tensor with $|\Omega|$ non-zero entries.

Furthermore, let $\{\Theta_1^*, \dots, \Theta_K^*, \boldsymbol{\mathcal{Z}}^*\}$ be the optimal solution of (5). The optimal solution of (3) is given by $\boldsymbol{\mathcal{W}}^* = \sum_k \boldsymbol{\mathcal{W}}^{(k)*}$, where $\boldsymbol{\mathcal{W}}^{(k)*} = \lambda_k(\boldsymbol{\mathcal{Z}}^* \times_k \Theta_k^*) \ \forall k$ and $\times_k$ denotes the tensor-matrix multiplication along mode k.

Proof. From using the auxiliary Θ_ks in (3), we obtain the following formulation:

$$\min_{\boldsymbol{\mathcal{W}}^{(k)}, k \in \{1, \dots, K\}} \left\| (\sum_{k=1}^{K} \boldsymbol{\mathcal{W}}^{(k)})_\Omega - \boldsymbol{\mathcal{Y}}_\Omega \right\|_F^2 + \sum_{k=1}^{K} \frac{1}{2\lambda_k} \langle \Theta_k^\dagger, \boldsymbol{\mathcal{W}}_k^{(k)} \boldsymbol{\mathcal{W}}_k^{(k)\top} \rangle. \tag{6}$$

For deriving the dual, we introduce new variables $\mathbf{A}_k, k \in \{1, \dots, K\}$ which satisfy the constraints $\mathbf{A}_k = \boldsymbol{\mathcal{W}}_k^{(k)}$. We now introduce the dual variables $\Lambda_k \in \mathbb{R}^{n_k \times n_1 \dots n_K / n_k}$ corresponding to those

additional constraints. The Lagrangian L of (6) is given as:

$$L(\boldsymbol{\mathcal{W}}^{(1)}, \ldots, \boldsymbol{\mathcal{W}}^{(K)}, \mathbf{A}_1, \ldots, \mathbf{A}_K, \Lambda_1, \ldots, \Lambda_K) = \left\| (\sum_{k}^{K} \boldsymbol{\mathcal{W}}^{(k)})_\Omega - \boldsymbol{\mathcal{Y}}_\Omega \right\|_F^2 +$$

$$\sum_{k=1}^{K} \frac{1}{2\lambda_k} \langle \Theta_k^\dagger, \mathbf{A}_k \mathbf{A}_k^\top \rangle + \sum_{k=1}^{K} \langle \Lambda_k, (\boldsymbol{\mathcal{W}}_k^{(k)} - \mathbf{A}_k) \rangle. \quad (7)$$

The dual function of (6) is defined as:

$$\min_{\substack{\boldsymbol{\mathcal{W}}^{(k)} \in \mathbb{R}^{n_1 \times \ldots \times n_K}, \\ \mathbf{A}_k \in \mathbb{R}^{n_k \times n_1 \cdots n_K / n_k}, \\ k \in \{1, \ldots, K\}}} L(\boldsymbol{\mathcal{W}}^{(1)}, \ldots, \boldsymbol{\mathcal{W}}^{(K)}, \mathbf{A}_1, \ldots, \mathbf{A}_K, \Lambda_1, \ldots, \Lambda_K). \quad (8)$$

By minimizing L with respect to $\boldsymbol{\mathcal{W}}^{(k)}$ and $\mathbf{A}_k$, for all $k \in \{1, \ldots, K\}$, we arrive at the following conditions:

$$[\Lambda_k]_k = 2(\boldsymbol{\mathcal{Y}}_\Omega - (\sum_{k=1}^{K} \boldsymbol{\mathcal{W}}^{(k)})_\Omega), \quad (9)$$

$$\mathbf{A}_k = \lambda_k \Theta_k \Lambda_k, \quad (10)$$

where $[\mathbf{P}]_k$ represents mode-k folding of a matrix $\mathbf{P}$ into a tensor.

From (9), it can be seen that all $[\Lambda_k]_k$ are equal, which we represent with $\boldsymbol{\mathcal{Z}}$. Therefore (9) and (10) can written as:

$$\boldsymbol{\mathcal{Z}} = 2(\boldsymbol{\mathcal{Y}}_\Omega - (\sum_{k=1}^{K} \boldsymbol{\mathcal{W}}^{(k)})_\Omega), \quad (11)$$

$$\mathbf{A}_k = \lambda_k \boldsymbol{\mathcal{Z}} \times_k \Theta_k, \quad (12)$$

From (11), it is clear that $\boldsymbol{\mathcal{Z}}_{(i_1, \ldots, i_K)}$ is non-zero only for $(i_1, \ldots, i_K) \in \Omega$, i.e., $\boldsymbol{\mathcal{Z}}$ is a sparse tensor, thereby ensuring the constraint $\boldsymbol{\mathcal{Z}} \in \mathcal{C}$ is satisfied. Using (11) and (12) in (8) gives the dual formulation and hence proving the theorem. □

Remark 1. *Theorem 1 shows that the optimal solutions $\boldsymbol{\mathcal{W}}^{(k)*}$ for all k of (3) are completely characterized by a single sparse tensor $\boldsymbol{\mathcal{Z}}^*$ and K low-rank positive semi-definite matrices $\{\Theta_1^*, \ldots, \Theta_K^*\}$. It should be noted that such a novel relationship of $\boldsymbol{\mathcal{W}}^{(k)*}$ (for all k) with each other is not evident from the formulation (3).*

We next present the following result related to the form of the optimal solution of (3).

Corollary 1. *(Representer theorem) The optimal solution of the primal problem (3) admits a representation of the form: $\boldsymbol{\mathcal{W}}^{(k)*} = \lambda_k(\boldsymbol{\mathcal{Z}} \times_k \Theta_k) \; \forall k$, where $\boldsymbol{\mathcal{Z}} \in \mathcal{C}$ and $\Theta_k \in \mathcal{P}^{n_k}$.*

Instead of solving the minimax problem (5) directly, we solve the following equivalent min optimization problem:

$$\min_{\theta \in \mathcal{P}^{n_1} \times \ldots \times \mathcal{P}^{n_K}} f(\theta), \quad (13)$$

where $\theta = (\Theta_1, \ldots, \Theta_K)$ and $f : \mathcal{P}^{n_1} \times \ldots \times \mathcal{P}^{n_K} \to \mathbb{R}$ is the convex function

$$f(\theta) := \max_{\boldsymbol{\mathcal{Z}} \in \mathcal{C}} \langle \boldsymbol{\mathcal{Z}}, \boldsymbol{\mathcal{Y}}_\Omega \rangle - \frac{1}{4} \|\boldsymbol{\mathcal{Z}}\|_F^2 - \sum_k \frac{\lambda_k}{2} \langle \Theta_k, \boldsymbol{\mathcal{Z}}_k \boldsymbol{\mathcal{Z}}_k^\top \rangle. \quad (14)$$

The problem formulation (13) allows to decouple the structures of the min and max problems in (5).

As discussed earlier in the section, the optimal $\Theta_k^* \in \mathcal{P}^{n_k}$ is a low-rank positive semi-definite matrix for all k. In spite of the low-rankness of the optimal solution, an algorithm for (5) need not produce intermediate iterates that are low rank. From the perspective of large-scale applications, this observation as well as other computational efficiency concerns discussed below motivate to exploit a fixed-rank parameterization of Θ_k for all k.

Fixed-rank parameterization of Θ_k

We propose to explicitly constrain the rank of Θ_k to r_k as follows:

$$\Theta_k = \mathbf{U}_k \mathbf{U}_k^\top, \tag{15}$$

where $\mathbf{U}_k \in \mathcal{S}_{r_k}^{n_k}$ and $\mathcal{S}_r^n := \{\mathbf{U} \in \mathbb{R}^{n \times r} : \|\mathbf{U}\|_F = 1\}$. In large-scale tensor completion problems, it is common to set $r_k \ll n_k$, where the fixed-rank parameterization (15) of Θ_k has a two-fold advantage. First, the search space dimension of (13) drastically reduces from $n_k((n_k + 1)/2 - 1)$, which is *quadratic* in tensor dimensions, to $n_k r_k - 1 - r_k(r_k - 1)/2$, which is *linear* in tensor dimensions [JBAS10]. Second, enforcing the constraint $\mathbf{U}_k \in \mathcal{S}_{r_k}^{n_k}$ costs $O(n_k r_k)$, which is *linear* in tensor dimensions and is computationally much cheaper than enforcing $\Theta_k \in \mathcal{P}^{n_k}$ that costs $O(n_k^3)$.

Fixed-rank dual formulation

A first formulation is obtained by employing the parameterization (15) directly in the problems (13) and (14). We subsequently solve the resulting problem as a minimization problem as follows

$$\text{Problem} \quad \mathcal{D}: \quad \min_{u \in \mathcal{S}_{r_1}^{n_1} \times \ldots \times \mathcal{S}_{r_K}^{n_K}} \quad g(u), \tag{16}$$

where $u = (\mathbf{U}_1, \ldots, \mathbf{U}_K)$ and $g : \mathcal{S}_{r_1}^{n_1} \times \ldots \times \mathcal{S}_{r_K}^{n_K} \to \mathbb{R}$ is the function

$$g(u) := \max_{\mathcal{Z} \in \mathcal{C}} \langle \mathcal{Z}, \mathcal{Y}_\Omega \rangle - \frac{1}{4} \|\mathcal{Z}\|_F^2 - \sum_k \frac{\lambda_k}{2} \left\| \mathbf{U}_k^\top \mathcal{Z}_k \right\|_F^2. \tag{17}$$

It should be noted that though (16) is a non-convex problem in u, the optimization problem in (17) is *strongly* convex in $\mathcal{Z}$ for a given u and has *unique* solution.

Fixed-rank least-squares formulation

The second formulation is motivated by the representer theorem (Corollary 1) and the fixed-rank parameterization (15). However, instead of solving (3), we take a more practical approach of solving the following low-rank

$$\text{Problem} \quad \mathcal{P}: \quad \min_{\substack{u \in \mathcal{S}_{r_1}^{n_1} \times \ldots \times \mathcal{S}_{r_K}^{n_K} \\ \mathcal{Z} \in \mathcal{C}}} \quad h(u, \mathcal{Z}) \tag{18}$$

and $h : \mathcal{S}_{r_1}^{n_1} \times \ldots \times \mathcal{S}_{r_K}^{n_K} \times \mathcal{C} \to \mathbb{R}$ is

$$h(u, \mathcal{Z}) := \left\| \left(\sum_k \lambda_k (\mathcal{Z} \times_k \mathbf{U}_k \mathbf{U}_k^\top) \right)_\Omega - \mathcal{Y}_\Omega \right\|_F^2. \tag{19}$$

where $\mathcal{W} = \sum_k \mathcal{W}^{(k)}$ and $\mathcal{W}^{(k)} = \lambda_k \mathcal{Z} \times_k \mathbf{U}_k \mathbf{U}_k^\top$. It should be noted that the objective function in (18) does not have an explicit regularizer as in (3) to ensure non-sparse $\mathcal{W}^{(k)}$ and low-rank $\mathcal{W}_k^{(k)}$. However, the regularizer is *implicit* since we employed the representer theorem and the fixed-rank parameterization ($\mathcal{Z}$ is common for all $\mathcal{W}^{(k)}$ and $\mathbf{U}_k \in \mathcal{S}_{r_k}^{n_k}$).
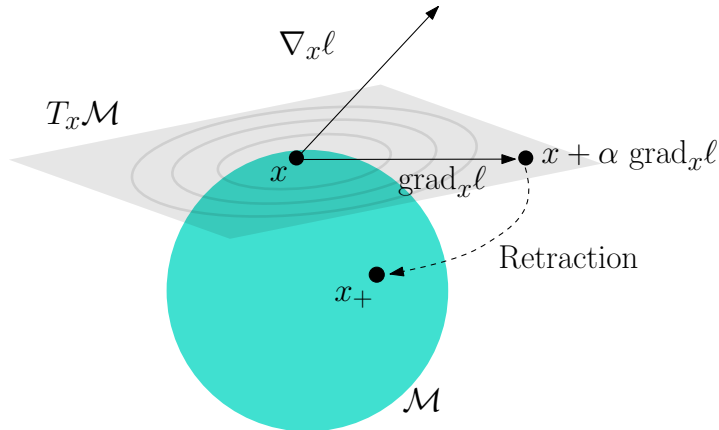
Figure 2: A schematic view of the gradient descent algorithm for minimizing a function ℓ on a manifold $\mathcal{M}$. The *linearization* of $\mathcal{M}$ at x is characterized by the tangent space $T_x\mathcal{M}$. The Riemannian gradient $\operatorname{grad}_x \ell$ is the steepest-descent direction of ℓ (in the tangent space $T_x\mathcal{M}$), which is derived from the standard gradient $\nabla_x\ell(x)$ (which need not lie in $T_x\mathcal{M}$). Following $\operatorname{grad}_x \ell$ with step size α leaves the manifold that is subsequently pulled back onto the manifold with the retraction operation to ensure strict feasibility. For the trust-region algorithm, additional care is taken to exploit the second-order geometry of the manifold.

4   Optimization algorithms

In this section we discuss the optimization algorithms for $\boldsymbol{\mathcal{D}}$ (16) and $\boldsymbol{\mathcal{P}}$ (18). Both of these are of the type

$$\min_{x \in \mathcal{M}} \quad \ell(x), \tag{20}$$

where $\ell : \mathcal{M} \to \mathbb{R}$ is a smooth loss and $\mathcal{M}$ is the constraint set. For Problem $\boldsymbol{\mathcal{D}}$, $\mathcal{M} := \mathcal{S}_{r_1}^{n_1} \times \ldots \times \mathcal{S}_{r_K}^{n_K}$ and Problem $\boldsymbol{\mathcal{P}}$, $\mathcal{M} := \mathcal{S}_{r_1}^{n_1} \times \ldots \times \mathcal{S}_{r_K}^{n_K} \times \mathcal{C}$.

In order to propose numerically efficient algorithms for optimization over $\mathcal{M}$, we exploit the particular structure of the set $\mathcal{S}_r^n$, which is known as the *spectrahedron* manifold [JBAS10]. The spectrahedron manifold has the structure of a compact Riemannian quotient manifold [JBAS10]. Consequently, optimization on the spectrahedron manifold is handled in the Riemannian optimization framework. This allows to exploit the rotational invariance of the constraint $\|\mathbf{U}\|_F = 1$ naturally. The Riemannian manifold optimization framework embeds the constraint $\|\mathbf{U}\|_F = 1$ into the search space, thereby translating the constrained optimization problem into *unconstrained* optimization problem over the spectrahedron manifold. The Riemannian framework generalizes a number of classical first- and second-order (e.g., the conjugate gradient and trust-region algorithms) Euclidean algorithms to manifolds and provides concrete convergence guarantees [EAS98, AMS08, SKM17, ZRS16, SI13]. The work [AMS08] in particular shows a systematic way of implementing trust-region (TR) algorithms on quotient manifolds. A full list of optimization-related ingredients and their matrix characterizations for the spectrahedron manifold $\mathcal{S}_r^n$ is in Section B. Overall, the constraint $\mathcal{M}$ is endowed a Riemannian structure. Figure 2 depicts a general overview of a gradient-based algorithm on $\mathcal{M}$.

We implement the Riemannian TR (second-order) algorithm for (20). To this end, we require the notions of the *Riemannian gradient* (the first-order derivative of the objective function on the manifold), the *Riemannian Hessian* along a search direction (the *covariant* derivative of the Riemannian gradient along a tangential direction on the manifold), and the *retraction* operator which ensures that we always stay on the manifold (i.e., maintain strict feasibility). The Riemannian

Algorithm 1 Proposed Riemannian trust-region algorithm for (16) and (18).

Input: $\boldsymbol{\mathcal{Y}}_\Omega$, rank $(r_1, \ldots, r_K)$, regularization parameter λ, and tolerance ϵ.
Initialize : $u \in \mathcal{M}$.
repeat
    1: Compute the gradient $\nabla_u \ell$ as given in Lemma 1 for Problem $\boldsymbol{\mathcal{D}}$ and Lemma 2 for Problem $\boldsymbol{\mathcal{P}}$.
    2: Compute the search direction which minimizes the trust-region subproblem. It makes use
       of $\nabla_u \ell$ and its directional derivative presented in Lemma 1 and Lemma 2 for $\boldsymbol{\mathcal{D}}$ and $\boldsymbol{\mathcal{P}}$,
       respectively.
    3: Update u with the retraction step to maintain strict feasibility on $\mathcal{M}$.
       Specifically for the spectrahedron manifold, $\mathbf{U}_k \leftarrow (\mathbf{U}_k + \mathbf{V}_k)/\|\mathbf{U}_k + \mathbf{V}_k\|_F$,
       where $\mathbf{V}_k$ is the search direction.
until $\|\nabla_u \ell\|_F < \epsilon$.
Output: u^*

---

gradient and Hessian notions require computations of the standard (Euclidean) gradient $\nabla_x \ell(x)$ and the directional derivative of this gradient along a given search direction v denoted by $\mathrm{D}\nabla_x \ell(x)[v]$, the expressions of which for $\boldsymbol{\mathcal{D}}$ (16) and $\boldsymbol{\mathcal{P}}$(18) are shown in Lemma 1 and Lemma 2, respectively.

Lemma 1. *Let* $\hat{\boldsymbol{\mathcal{Z}}}$ *be the optimal solution of the convex problem (17) at* $u \in \mathcal{S}_{r_1}^{n_1} \times \ldots \times \mathcal{S}_{r_K}^{n_K}$. *Let* $\nabla_u g$ *denote the gradient of* $g(u)$ *at* u *and* $\mathrm{D}\nabla_u g[v]$ *denote the directional derivative of the gradient* $\nabla_u g$ *along* $v \in \mathbb{R}^{n_1 \times r_1} \times \ldots \times \mathbb{R}^{n_K \times r_K}$. *Let* $\dot{\boldsymbol{\mathcal{Z}}}_k$ *denote the directional derivative of* $\boldsymbol{\mathcal{Z}}_k$ *along* v *at* $\hat{\boldsymbol{\mathcal{Z}}}_k$. *Then,*

$$\begin{aligned} \nabla_u g &= (-\lambda_1 \hat{\boldsymbol{\mathcal{Z}}}_1 \hat{\boldsymbol{\mathcal{Z}}}_1^\top \mathbf{U}_1, \ldots, -\lambda_K \hat{\boldsymbol{\mathcal{Z}}}_K \hat{\boldsymbol{\mathcal{Z}}}_K^\top \mathbf{U}_K) \\ \mathrm{D}\nabla_u g[v] &= (-\lambda_1 \mathbf{A}_1, \ldots, -\lambda_K \mathbf{A}_K), \end{aligned}$$

where $\mathbf{A}_k = \hat{\boldsymbol{\mathcal{Z}}}_k \hat{\boldsymbol{\mathcal{Z}}}_k^\top \mathbf{V}_k + \mathrm{symm}(\dot{\boldsymbol{\mathcal{Z}}}_k \hat{\boldsymbol{\mathcal{Z}}}_k^\top) \mathbf{U}_k$ *and* symm *extracts the symmetric factor of a matrix, i.e.,* $\mathrm{symm}(\boldsymbol{\Delta}) = (\boldsymbol{\Delta} + \boldsymbol{\Delta}^\top)/2$.

Proof. The gradient is computed by employing the Danskin's theorem [Ber99, BS00]. The directional derivative of the gradient with respect to $\mathbf{U}$ follows directly from the chain rule keeping $\boldsymbol{\mathcal{Z}}$ fixed (to the optimal solution of (17) for given $\mathbf{U}$). $\qquad\square$

A key requirement in Lemma 1 is solving (17) for $\hat{\boldsymbol{\mathcal{Z}}}$ computationally efficiently for a given $u = (\mathbf{U}_1, \ldots, \mathbf{U}_K)$. It should be noted that (17) has a closed-form sparse solution, which is equivalent to solving the *linear* system

$$\hat{\boldsymbol{\mathcal{Z}}}_\Omega + \sum_k \lambda_k (\hat{\boldsymbol{\mathcal{Z}}}_\Omega \times_k \mathbf{U}_k \mathbf{U}_k^\top)_\Omega = \boldsymbol{\mathcal{Y}}_\Omega. \tag{21}$$

Solving the linear system (21) in a *single* step is computationally expensive (it involves the use of Kronecker products, vectorization of a sparse tensor, and a matrix inversion). Instead, we use an *iterative* solver that exploits the sparsity in the variable $\boldsymbol{\mathcal{Z}}$ and the factorization form $\mathbf{U}_k \mathbf{U}_k^\top$ efficiently. Similarly, given $\hat{\boldsymbol{\mathcal{Z}}}$ and v, $\dot{\boldsymbol{\mathcal{Z}}}$ can be computed by solving

$$\dot{\boldsymbol{\mathcal{Z}}}_\Omega + \sum_k \lambda_k (\dot{\boldsymbol{\mathcal{Z}}}_\Omega \times_k \mathbf{U}_k \mathbf{U}_k^\top)_\Omega = -\sum_k \lambda_k (\hat{\boldsymbol{\mathcal{Z}}}_\Omega \times_k (\mathbf{V}_k \mathbf{U}_k^\top + \mathbf{U}_k \mathbf{V}_k^\top))_\Omega. \tag{22}$$

Lemma 2. *The gradient of* $h(u, \boldsymbol{\mathcal{Z}})$ *in (19) at* $(u, \boldsymbol{\mathcal{Z}})$ *and their directional derivatives along* $(v, \dot{\boldsymbol{\mathcal{Z}}})$, *where* $v \in \mathbb{R}^{n_1 \times r_1} \times \ldots \times \mathbb{R}^{n_K \times r_K}$ *and* $\dot{\boldsymbol{\mathcal{Z}}} \in \mathcal{C}$ *are given as follows.*

$$\begin{aligned} \nabla_u h &= (\lambda_1 \mathrm{symm}(\boldsymbol{\mathcal{R}}_1 \boldsymbol{\mathcal{Z}}_1^\top) \mathbf{U}_1, \ldots, \lambda_K \mathrm{symm}(\boldsymbol{\mathcal{R}}_K \boldsymbol{\mathcal{Z}}_K^\top) \mathbf{U}_K) \\ \nabla_{\boldsymbol{\mathcal{Z}}} h &= \sum_k \lambda_k (\boldsymbol{\mathcal{R}} \times_k \mathbf{U}_k \mathbf{U}_k^\top)_\Omega \\ \mathrm{D}\nabla_u h[(v, \dot{\boldsymbol{\mathcal{Z}}})] &= (\lambda_1 \mathbf{B}_1, \ldots, \lambda_K \mathbf{B}_K) \\ \mathrm{D}\nabla_{\boldsymbol{\mathcal{Z}}} h[(v, \dot{\boldsymbol{\mathcal{Z}}})] &= \sum_k \lambda_k (\dot{\boldsymbol{\mathcal{R}}} \times_k \mathbf{U}_k \mathbf{U}_k^\top + \boldsymbol{\mathcal{R}} \times_k \mathrm{symm}(\mathbf{U}_k \mathbf{V}_k^\top))_\Omega, \end{aligned}$$

where

$$\mathcal{R} = \sum_k \lambda_k (\boldsymbol{\mathcal{Z}} \times_k \mathbf{U}_k \mathbf{U}_k^\top)_\Omega - \boldsymbol{\mathcal{Y}}_\Omega$$
$$\mathbf{B}_k = \mathrm{symm}(\dot{\mathcal{R}}_k \boldsymbol{\mathcal{Z}}_k^\top + \mathcal{R}_k \dot{\boldsymbol{\mathcal{Z}}}_k^\top)\mathbf{U}_k + \mathrm{symm}(\mathcal{R}_k \boldsymbol{\mathcal{Z}}_k^\top)\mathbf{V}_k \qquad (23)$$
$$\dot{\mathcal{R}} = \sum_k \lambda_k (\dot{\boldsymbol{\mathcal{Z}}} \times_k \mathbf{U}_k \mathbf{U}_k^\top + \boldsymbol{\mathcal{Z}} \times_k \mathrm{symm}(\mathbf{U}_k \mathbf{V}_k^\top))_\Omega.$$

Proof. The directional derivatives of the gradient with respect to $\mathbf{U}$ and $\boldsymbol{\mathcal{Z}}$ follow directly from the chain rule. □

The Riemannian TR algorithm solves a Riemannian trust-region sub-problem in every iteration [AMS08, Chapter 7]. The TR sub-problem is a *second-order* approximation of the objective function in a neighborhood, solution to which does not require inverting the full Hessian of the objective function. It makes use of the gradient $\nabla_x \ell$ and its directional derivative along a search direction. The TR sub-problem is approximately solved with an iterative solver, e.g., the truncated conjugate gradient algorithm. The TR sub-problem outputs a potential update candidate for x, which is then accepted or rejected based on the amount of decrease in ℓ. Algorithm 1 summarizes the key steps of the TR algorithm for solving (20).

Remark 2: In contrast to Problem $\boldsymbol{\mathcal{D}}$, where (17) is required to be solved for $\boldsymbol{\mathcal{Z}}$ at a given u, in Problem $\boldsymbol{\mathcal{P}}$ both u and $\boldsymbol{\mathcal{Z}}$ are updated simultaneously.

Computational and memory complexity

The per-iteration computational complexities of our Riemannian TR algorithms for both the formulations – $\boldsymbol{\mathcal{D}}$ (16) and $\boldsymbol{\mathcal{P}}$ (18) – are shown in Table 1. The computational complexity scales *linearly* with the number of known entries $\boldsymbol{\mathcal{Y}}_\Omega$, denoted by $|\Omega|$.

In particular, the per-iteration computational cost depends on two sets of operations. First, on manifold related ingredients like the retraction operation. Second, the objective function related ingredients like the computation of the partial derivatives.

The manifold related operations cost $O(\sum_k n_k r_k^2 + r_k^3)$ for $\boldsymbol{\mathcal{D}}$ (16) and $O(|\Omega| + \sum_k n_k r_k^2 + r_k^3)$ for $\boldsymbol{\mathcal{P}}$ (18). Specifically, the retraction on the spectrahedron manifold $\mathcal{S}_{r_k}^{n_k}$ costs $O(n_k r_k)$ as it needs to project a matrix of size $n_k \times r_k$ on to the set $\mathcal{S}_{r_k}^{n_k}$, which costs $O(n_k r_k)$.

The following are the computational cost for the objective function related ingredients.

- $\mathbf{U}_k^\top \boldsymbol{\mathcal{Z}}_k$**.** It involves computation of $n_k \times r_k$ matrix $\mathbf{U}_k$ with mode-k unfolding of a sparse $\boldsymbol{\mathcal{Z}}$ with $|\Omega|$ non-zero entries. Such matrix-matrix operations are used in both $\boldsymbol{\mathcal{D}}$ (16) and $\boldsymbol{\mathcal{P}}$ (18). The computation of $\mathbf{U}_k^\top \boldsymbol{\mathcal{Z}}_k$ costs $O(|\Omega| r_k)$. It should be noted that although the dimension of $\boldsymbol{\mathcal{Z}}_k$ is $n_k \times \prod_{i=1, i\neq k}^K n_i$, only a maximum of $|\Omega|$ columns have non-zero entries. We exploit this property of $\boldsymbol{\mathcal{Z}}_k$ and have a compact memory storage of $\mathbf{U}_k^\top \boldsymbol{\mathcal{Z}}_k$.

- **Computation of the solution $\hat{\boldsymbol{\mathcal{Z}}}$ and $\dot{\boldsymbol{\mathcal{Z}}}$ used in $\boldsymbol{\mathcal{D}}$ (16).** We use an iterative solver for (21) and (22), which requires computing the left hand side of (21) for a given candidate $\boldsymbol{\mathcal{Z}}$. If m is the number of iterations allowed, then the computational cost is costs $O(m|\Omega| \sum_k r_k)$.

- **Computation of $g(u)$ in $\boldsymbol{\mathcal{D}}$ (16) and its partial derivatives.** The computation of $g(u)$ relies on the solution of (21) and then explicitly computing the objective function in (17). This costs $O(m|\Omega| \sum_k r_k + K|\Omega|)$, where m is the number of iterations to solve (21). The computation of $\nabla_u g$ requires the computation of terms like $\hat{\boldsymbol{\mathcal{Z}}}_k(\hat{\boldsymbol{\mathcal{Z}}}_k^\top \mathbf{U}_k)$, and overall it costs $O(|\Omega| \sum_k r_k)$. Given a search direction v, the computation of $\mathrm{D}\nabla_u g[v]$ costs $O(|\Omega| \sum_k r_k)$.

- **Computation of $h(u, \boldsymbol{\mathcal{Z}})$ in $\boldsymbol{\mathcal{P}}$ (18) and its partial derivatives.** The computation of $h(u, \boldsymbol{\mathcal{Z}})$ costs $O(\sum_k |\Omega| r_k)$. Specifically, the computation of $\mathcal{R}$ costs $O(|\Omega| \sum_k r_k)$. The computations of $\nabla_u h$ and it directional derivatives require the computation of terms like

Table 1: The per-iteration computational complexity for the TR algorithms for $\boldsymbol{\mathcal{D}}$ (16) and $\boldsymbol{\mathcal{P}}$(18). For $\boldsymbol{\mathcal{D}}$, m is the number of iterations needed to solve (21) and (22) approximately.

Formulation	Computational cost
$\boldsymbol{\mathcal{D}}$ (16)	$O(m\|\Omega\|\sum_k r_k + \sum_k n_k r_k^2 + \sum_k r_k^3)$
$\boldsymbol{\mathcal{P}}$ (18)	$O(\|\Omega\|\sum_k r_k + \sum_k n_k r_k^2 + \sum_k r_k^3)$

$\mathrm{symm}(\boldsymbol{\mathcal{R}}_k \boldsymbol{\mathcal{Z}}_k^\top)\mathbf{U}_k$, which costs $O(|\Omega|\sum_k r_k)$. The same computational overhead is for $\nabla_{\boldsymbol{\mathcal{Z}}} h$ and its directional derivative.

The memory cost for both algorithms is $O(|\Omega| + \sum_k n_k r_k)$, where $|\Omega|$ is linear in n_k.

Convergence

The Riemannian TR algorithms come with rigorous convergence guarantees. [AMS08] discusses the rate of convergence analysis of manifold algorithms, which directly apply in our case. Specifically for trust regions, the global convergence to a first-order critical point is discussed in [AMS08, Section 7.4.1] and the local convergence to local minima is discussed in [AMS08, Section 7.4.2]. From an implementation perspective, we follow the existing approaches [KSV14, KM16, GYK17] and upper-limit the number of TR iterations.

Numerical implementation

Our algorithms are implemented in the Manopt toolbox [BMAS14] in Matlab, which has off-the-shelf generic TR implementation. For efficient computation of certain sparse matrix-vector we make use of the mex support in Matlab. For solving (21) and (22), which are used in solving $\boldsymbol{\mathcal{D}}$ (16), we rely on the conjugate gradients solver of Matlab.

5   Experiments

In this section, we evaluate the generalization performance and efficiency of the proposed algorithms against state-of-the-art in several tensor completion applications. We provide the experimental analysis of our proposed algorithms for the applications of video and hyperspectral-image completion, recommendation and link prediction. Table 2 provides the details of all the datasets we experiment with. We compare performance of the following trace norm based algorithms.

1. TR-MM: our trust-region algorithm for the proposed formulation $\boldsymbol{\mathcal{D}}$ (16).

2. TR-LS: our trust-region algorithm for the proposed formulation $\boldsymbol{\mathcal{P}}$ (18).

3. Latent[1] [THK10]: a convex optimization based algorithm with latent trace norm.

4. Hard [SDLS14]: an algorithm based on a convex optimization framework with spectral regularization for tensor learning.

5. HaLRTC[2] [LMWY13]: an ADMM based algorithm for tensor completion.

6. FFW[3] [GYK17]: a Frank-Wolfe algorithm with scaled latent trace norm regularization.

---

[1] http://tomioka.dk/softwares/
[2] http://www.cs.rochester.edu/u/jliu/code/TensorCompletion.zip
[3] http://home.cse.ust.hk/~qyaoaa/

Table 2: Summary of the datasets.

Dataset	Dimensions	Task
Ribeira	$203 \times 268 \times 33$	Image completion
Tomato	$242 \times 320 \times 167$	Video completion
Baboon	$256 \times 256 \times 3$	Image completion
FB15k-237	$14541 \times 14541 \times 237$	Link prediction
YouTube (subset)	$1509 \times 1509 \times 5$	Link prediction
YouTube (full)	$15088 \times 15088 \times 5$	Link prediction
ML10M	$71567 \times 10681 \times 731$	Recommendation

Table 3: Summary of baseline algorithms.

Algorithm	Modeling details
	Trace norm based algorithms
FFW	Scaled latent trace norm (scaled $S_1/1$-norm) + Frank Wolfe optimization + basis size reduction
Hard	Scaled overlapped trace norm + proximal gradient
HaLRTC	Scaled overlapped trace norm + Alternating direction methods of multipliers (ADMM)
Latent	Latent trace norm ($S_1/1$-norm) + ADMM
	Other algorithms
Rprecon	Fixed multilinear rank + Riemannian CG with preconditioning
BayesCP	Bayesian CP algorithm with rank tuning
geomCG	Riemannian fixed multilinear rank CG algorithm
Topt	Fixed multilinear rank CG algorithm
T-svd	Tensor tubal-rank + ADMM

We denote our algorithms for Problem $\mathcal{D}$ (16) and Problem $\mathcal{P}$ (18) as TR-MM and TR-LS, respectively. For both of our algorithms, we set $\lambda_k = \lambda n_k$ for all k, which implies that in $\lambda > 0$ is the only hyper-parameter that needed to be tuned. We tune λ with 5-fold cross-validation on the train data of the split from the range $\{10^{-3}, 10^{-2}, \ldots, 10^3\}$.

Latent trace norm based algorithms such as FFW [GYK17] and Latent [THK10], and overlapped trace norm based algorithms such as HaLRTC [LMWY13] and Hard [SDLS14] are the closest to our approach. FFW is a recently proposed state-of-the-art large scale tensor completion algorithm. Table 3 summarizes the baseline algorithms.

The proposed algorithms are implemented using Manopt toolbox [BMAS14] in Matlab. All the experiments are run on an Intel Xeon CPU with 32GB RAM and 2.40GHz processor.

5.1   Video and hyperspectral-image completion

We work with the following data sets for predicting missing values in multi-media data like videos and hyperspectral images.

(a) Ribeira　　　　　　　　(b) Baboon

Figure 3: Image datasets used in our work.

Table 4: Mean test RMSE (lower is better) for hyperspectral-image completion, video completion, and recommendation problems. Our algorithms, TR-MM and TR-LS, obtain the best performance among trace norm based algorithms. '-' denotes the data set is too large for the algorithm to generate result.

	Ribeira	Tomato	Baboon	MovieLens10M
TR-MM	0.067 ± 0.001	$\mathbf{0.041 \pm 0.001}$	$\mathbf{0.121 \pm 0.001}$	0.840 ± 0.001
TR-LS	$\mathbf{0.064 \pm 0.002}$	0.047 ± 0.001	0.129 ± 0.002	0.838 ± 0.001
FFW	0.088 ± 0.001	0.045 ± 0.001	0.133 ± 0.001	0.895 ± 0.001
Rprecon	0.083 ± 0.001	0.052 ± 0.000	0.128 ± 0.001	$\mathbf{0.831 \pm 0.002}$
geomCG	0.156 ± 0.023	0.052 ± 0.000	0.128 ± 0.001	0.844 ± 0.003
Hard	0.114 ± 0.001	0.060 ± 0.001	0.126 ± 0.001	-
Topt	0.127 ± 0.023	0.102 ± 0.001	0.130 ± 0.004	-
HaLRTC	0.095 ± 0.001	0.202 ± 0.005	0.247 ± 0.012	-
Latent	0.087 ± 0.001	0.046 ± 0.001	0.459 ± 0.002	-
T-svd	0.064 ± 0.001	0.042 ± 0.001	0.146 ± 0.002	-
BayesCP	0.154 ± 0.001	0.103 ± 0.001	0.159 ± 0.001	-

Ribeira: A hyperspectral image data set[4] [FNA04] of size $1017 \times 1340 \times 33$, where each slice represents a particular image measured at a different wavelength. We re-size this tensor to $203 \times 268 \times 33$ [SDLS14, KSV14, KM16]. Following [KM16], the incomplete tensor as training data is generated by randomly sampling 10% of the entries and the evaluation (testing) of the learned tensor was done on another 10% of the entries. The experimented is repeated 10 times. Figure 3(a) shows the Ribeira image.

Tomato: A tomato video sequence data set[5] [LMWY13, CC14] of size $242 \times 320 \times 167$. We follow the same experimental setup as for Ribeira data set discussed above.

Baboon: An RGB image [6] of a baboon used in [ZZC15] which is modeled as a $256 \times 256 \times 3$ tensor. The experimental setup is same as that for Ribeira. Figure 3(b) shows the image.

(a) Ribeira

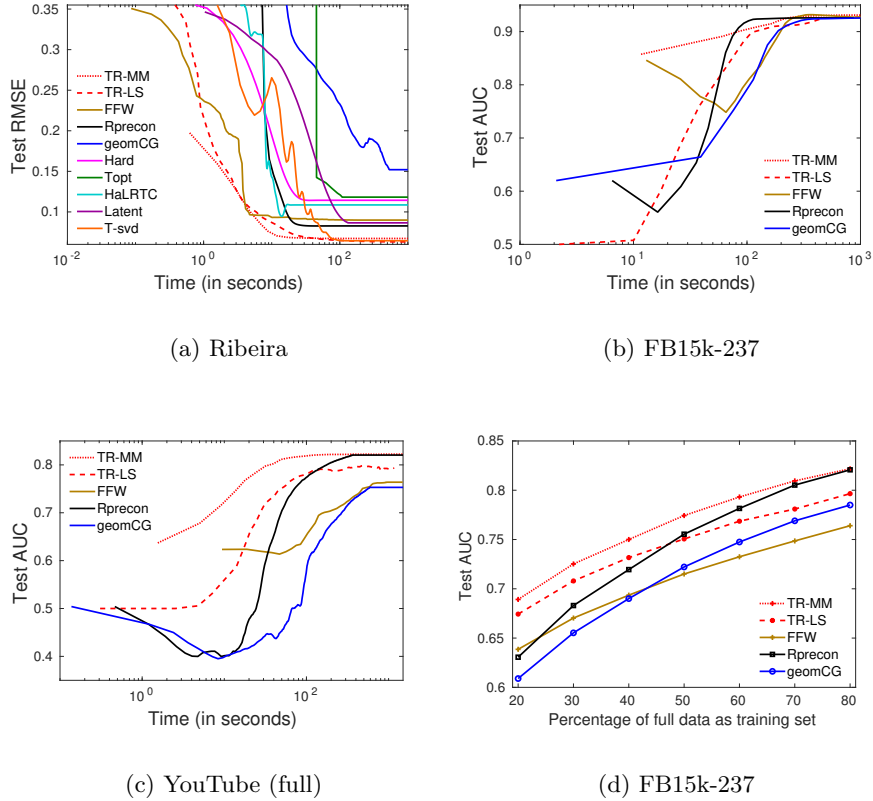(b) FB15k-237

(c) YouTube (full)

(d) FB15k-237

Figure 4: (a) Evolution of test RMSE on Ribeira; (b) & (c) Evolution of test AUC on FB15k-237 and YouTube, respectively. Both our algorithms, TR-MM and TR-LS, are the fastest to obtain a good generalization performance in all the three data sets; (d) Variation of test AUC as the amount of training data changes on FB15k-237. Our algorithms perform significantly better than others when the amount of training data is less.

Results. Table 4 reports the root mean squared error (RMSE) on the test set that is averaged over ten splits. Our algorithms, TR-MM and TR-LS, obtain the best results, outperforming other trace norm based algorithms on Ribeira, Tomato and Baboon data sets. Figure 4(a) shows the trade-off between the test RMSE and the training time of all the algorithms on Ribeira. It can be observed that both our algorithms converge to the lowest RMSE at a significantly faster rate compared to the other baselines. It is evident from the results that learning a mixture of non-sparse tensors, as learned by the proposed algorithms, helps in achieving better performance compared to the algorithms that learn a skewed sparse mixture of tensors.

5.2 Link prediction

In this application, the task is to predict missing or new links in knowledge graphs, social networks, etc. We consider the following data sets.
FB15k-237: FB15k-237[7] [TCP$^+$15] is a subset of FB15k dataset [BUGD$^+$13], containing facts of the form subject-predicate-object (RDF) triples from Freebase knowledge graph. The subject and object noun phrases are called entities and the predicates are the relationships among them. For

---

[7] http://kristinatoutanova.com/

instance, in the RDF triple {'Obama', 'president of', 'United States'}, 'Obama' and 'United States' are entities, and 'president of' is a relationship between them. The task is to predict relationships (from a given set of relations) between a pair of entities. FB15k-237 contains $14\,951$ entities and 237 relationships. It has $310\,116$ observed relationships (links) between pairs of entities, which are *positive* samples. In addition, $516\,606$ *negative* samples are generated following the procedure described in [BUGD+13]. Negative samples are those {entity$_a$, relationship$_c$, entity$_b$} triples in which the relationship$_c$ does not hold between entity$_a$ and entity$_b$. We model this task as a $14\,951 \times 14\,951 \times 237$ tensor completion problem. $\boldsymbol{\mathcal{Y}}_{(a,b,c)} = 1$ implies that relationship$_c$ exists between entity$_a$ and entity$_c$ (positive sample), and $\boldsymbol{\mathcal{Y}}_{(a,b,c)} = 0$ implies otherwise (negative sample). We keep 80% of the observed entries for training and the remaining 20% for testing.

YouTube: This is a link prediction data set[8] [TWL09] having 5 types of interactions between $15\,088$ users. The task is to predict the interaction (from a given set of interactions) between a pair of users. We model it as a $15\,088 \times 15\,088 \times 5$ tensor completion problem. All the entries are known in this case. We randomly sample 0.8% of the data for training [GYK17] and another 0.8% for testing.

It should be noted that Hard, HaLRTC, and Latent do not scale to the full FB15k-237 and YouTube data sets as they need to store full tensor in memory. Hence, we follow [GYK17] to create a subset of the YouTube data set of size $1509 \times 1509 \times 5$ in which 1509 users with most number of links are chosen. We randomly sample 5% of the data for training and another 5% for testing.

Results. All the above experiments are repeated on ten random train-test splits. Following [LZ11, GYK17], the generalization performance for link prediction task is measured by computing the area under the ROC curve on the test set (test AUC) for each algorithm. Table 5 report the average test AUC on YouTube (subset), Youtube (full) and FB15k-237 data sets. The TR-MM algorithm achieves the best performance in all the link prediction tasks, while TR-LS achieves competitive performance. This shows that the non-sparse mixture of tensors learned by TR-MM helps in achieving better performance. Figures 4 (b) & (c) plots the trade-off between the test AUC and the training time for FB15k-237 and YouTube, respectively. We can observe that TR-MM is the fastest to converge to a good AUC and take only a few iterations.

We also conduct experiments to evaluate the performance of different algorithms in challenging scenarios when the amount of training data available is less. On the FB15k-237 data set, we vary the size of training data from 20% to 80% of the observed entries, and the remaining 20% of the observed entries is kept as the test set. Figure 4 (d) plots the performance of different algorithms on this experiment. We can observe that both our algorithms do significantly better than baselines in data scarce regimes.

5.3 Movie recommendation

We also evaluate the algorithms on the MovieLens10M[9] data set [HK15]. This is a movie recommendation task — predict the ratings given to movies by various users. MovieLens10M contains $10\,000\,054$ ratings of $10\,681$ movies given by $71\,567$ users. Following [KM16], we split the time into 7-days wide bins, forming a tensor of size $71\,567 \times 10\,681 \times 731$. For our experiments, we generate 10 random train-test splits, where 80% of the observed entries is kept for training and remaining 20% is used for testing. Table 4 reports the average test RMSE on this task. It can be observed that the proposed algorithms outperform state-of-the-art latent and scaled latent trace norm based algorithms.

---

[8]http://leitang.net/data/youtube-data.tar
[9]http://files.grouplens.org/datasets/movielens/ml-10m-README.html

Table 5: Mean test AUC (higher is better) for link prediction problem. Our algorithms, TR-MM and TR-LS, obtain the best performance. '-' denotes the data set is too large for the algorithm to generate result.

	YouTube (subset)	YouTube (full)	FB15k-237
TR-MM	$\mathbf{0.957 \pm 0.001}$	$\mathbf{0.932 \pm 0.001}$	$\mathbf{0.823 \pm 0.001}$
TR-LS	0.954 ± 0.001	0.928 ± 0.002	0.799 ± 0.002
FFW	0.954 ± 0.001	0.929 ± 0.001	0.764 ± 0.003
Rprecon	0.941 ± 0.001	0.926 ± 0.001	0.821 ± 0.003
geomCG	0.941 ± 0.001	0.926 ± 0.001	0.785 ± 0.014
Hard	0.954 ± 0.001	-	-
Topt	0.941 ± 0.001	-	-
HaLRTC	0.783 ± 0.090	-	-
Latent	0.945 ± 0.001	-	-
T-svd	0.941 ± 0.001		
BayesCP	0.949 ± 0.002	-	-

Table 6: Ranks set for different datasets, at which the respective algorithms achieve best performance. For Tucker based algorithms, the rank is the Tucker rank which is different from the rank for the proposed algorithms.

Dataset	TR-MM/TR-LS	Tucker-based algorithms
Ribeira	(5,5,5)	(15,15,6)
Tomato	(10,10,10)	(15,15,15)
Baboon	(4,4,3)	(4,4,3)
FB15k-237	(20,20,1)	(5,5,5)
YouTube (subset)	(3,3,1)	(5,5,5)
YouTube (full)	(3,3,1)	(5,5,5)
ML10M	(20,10,1)	(4,4,4)

5.4   Results compared to other baseline algorithms

In addition to the trace norm based algorithms, we also compare against the following tensor decomposition algorithms:

1. Rprecon[10] [KM16]: a Riemannian manifold preconditioning algorithm for the fixed multi-linear rank low-rank tensor completion problem.

2. geomCG[11] [KSV14]: a fixed multi-linear rank low-rank tensor completion algorithm using Riemannian optimization on the manifold of fixed multi-linear rank tensors.

3. Topt [FJ15]: a non-linear conjugate gradient algorithm for Tucker decomposition.

4. T-svd[12][ZEA+14]: a tubal rank based tensor completion algorithm.

---

[10]https://bamdevmishra.com/codes/tensorcompletion/
[11]http://anchp.epfl.ch/geomCG
[12]http://www.ece.tufts.edu/~shuchin/software.html

Table 7: Results on outlier robustness experiments. Our algorithm, TR-MM, is more robust to outliers than the competing baselines. The symbol '$-$' denotes the dataset is too large for the algorithm to generate result.

	x	TR-MM	TR-LS	FFW	Rprecon	geomCG	Hard	Topt	HaLRTC	Latent	T-svd	BayesCP
Ribeira	0.05	**0.081**	0.137	0.095	0.157	0.258	0.142	0.169	0.121	0.103	0.146	0.201
(RMSE)	0.10	**0.111**	0.171	0.112	0.172	0.373	0.158	0.188	0.135	0.120	0.182	0.204
FB15k-237	0.05	**0.803**	0.768	0.734	0.794	0.764	$-$	$-$	$-$	$-$	$-$	$-$
(AUC)	0.10	**0.772**	0.736	0.711	0.765	0.739	$-$	$-$	$-$	$-$	$-$	$-$

5. BayesCP[13][ZZC15]:a Bayesian CP algorithm for tensor completion which has inbuilt rank tuning.

As can be observed from Tables 4 and 5, the proposed TR-MM and TR-LS achieve better results than the baseline algorithms. However for the movie recommendation task, Rprecon achieves better results. It should be noted that Topt, T-svd, and BayesCP are not scalable for large scale datasets.

For all the datasets, we vary the rank between (3,3,3) and (25,25,25) for TR-MM and TR-LS algorithms. For Tucker based algorithms we vary the Tucker rank between (3,3,3) and (25,25,25). Table 6 shows the ranks at which the algorithms achieve best results with different datasets.

Algorithms based on the Tucker decomposition [FJ15, KM16, KSV14] model the tensor completion problem as a factorization problem. Tucker decomposition is a form of higher order PCA [KB09], it decomposes a tensor into a core tensor multiplied by a matrix along each mode. Given the in complete tensor $\boldsymbol{\mathcal{Y}}_\Omega$, the Tucker decomposition based tensor completion problem is:

$$\min_{\substack{\boldsymbol{\mathcal{G}} \in \mathbb{R}^{r_1 \times \cdots \times r_K} \\ \mathbf{A}_k \in \mathbb{R}^{n_k \times r_k}, k \in \{1, \ldots, K\}}} \left\| \boldsymbol{\mathcal{Y}}_\Omega - (\boldsymbol{\mathcal{G}} \times_1 \mathbf{A}_1 \times_2 \ldots \times_K \mathbf{A}_K)_\Omega \right\|,$$

where the multi-linear rank of the optimal solution is given by $(r_1, r_2, \ldots, r_K)$, which is a user provided parameter in practical settings. CP based algorithms like BayesCP [ZZC15] are a special case of Tucker based modeling with $\boldsymbol{\mathcal{G}}$ set to be super diagonal and $r_1 = r_2 = \ldots = r_K = r$, which is a user provided parameter in practical setting. In particular, [ZZC15] discusses a probabilistic model. T-svd algorithm [ZEA$^+$14] models the tensor completion problem using the tensor singular value decomposition, and the rank here is the tensor tubal rank. Finally, algorithms based on latent norm regularization [GYK17, LMWY13, SDLS14, THK10] model the completion problem by approximating the input tensor as a sum of individual tensors (3). Each of these individual tensors is constrained to have a low-rank in one mode. The low-rank constraint is enforced on a different mode for each tensor. Due to the fundamental difference in modeling of these algorithms, the tensor ranks of the optimal solution obtained from algorithms that follow one approach cannot be compared with that of the algorithms that follow the other approach.

5.5 Results on outlier robustness

We also evaluate TR-MM and the baselines considered in Table 3 for outlier robustness on hyperspectal image completion and link prediction problems. In the Ribeira dataset, we add the standard Gaussian noise ($N(0, 1)$) to randomly selected x fraction of the entries in the training set. The minimum and the maximum value of entries in the (original) Ribeira are 0.01 and 2.09, respectively. In FB15k-237 dataset, we flip randomly selected x fraction of the entries in the training set, *i.e.*, the link is removed if present and vice-versa. We experiment with $x = 0.05$ and $x = 0.10$.

---

[13]https://github.com/qbzhao/BCPF

The results are reported in Table 7. We observe that our algorithm, TR-MM, obtains the best generalization performance and, hence, is the most robust to outliers. We also observe that trace norm regularized algorithms (TR-MM, FFW, Hard, HaLTRC, Latent) are relatively more robust to outliers than Tucker-decomposition based algorithms (Rprecon, geomCG, Topt), CP-decomposition based algorithm (BayesCP), and tensor tubal-rank based algorithm (T-svd).

6 Conclusion

In this paper, we introduce a variant of trace norm regularizer for low-rank tensor completion problem which learns the tensor as a non-sparse combination of tensors. The number of tensors in this combination is equal to the number of modes (K). Existing works [THK10, TS13, WST14, GYK17] learn a sparse combination of tensors, essentially learning the tensor as a low-rank matrix and losing higher order information in the available data. Hence, we recommend learning a non-sparse combination of tensors in trace norm regularized setting, especially since K is typically a small integer in most real-world applications. In our experiments, we observe better generalization performance with the proposed regularization. Theoretically, we provide the following result on the reconstruction error in the context of recovering an unknown tensor $\boldsymbol{\mathcal{W}}^*$ from noisy observation.

Lemma 3. *Let $\boldsymbol{\mathcal{W}}^*$ be the true tensor to be recovered from observed $\boldsymbol{\mathcal{Y}}$, which is obtained as $\boldsymbol{\mathcal{Y}} = \boldsymbol{\mathcal{W}}^* + \boldsymbol{\mathcal{E}}$, where $\boldsymbol{\mathcal{E}} \in \mathbb{R}^{n_1 \times \dots \times n_K}$ is the noise tensor. Assume that the regularization constant λ satisfies $\lambda \leq 1/(\sum_{k=1}^{K} \|\boldsymbol{\mathcal{E}}_k\|_\infty^2)^{1/2}$ then the estimator*

$$\hat{\boldsymbol{\mathcal{W}}} = \underset{\boldsymbol{\mathcal{W}}}{argmin}(\frac{1}{2}\|\boldsymbol{\mathcal{Y}} - \boldsymbol{\mathcal{W}}\|_F^2 + \frac{1}{\lambda}\sum_k \|\boldsymbol{\mathcal{W}}_k^{(k)}\|_*^2),$$

satisfies the inequality $\|\hat{\boldsymbol{\mathcal{W}}} - \boldsymbol{\mathcal{W}}^*\|_F \leq \frac{2}{\lambda}\sqrt{\min_k n_k}$. When noise approaches zero, i.e., $\boldsymbol{\mathcal{E}} \to 0$, the right hand side also approaches zero.

The proof of the above lemma is in Section A. A similar result on the *latent trace norm*, which learns a sparse combination of tensors, is presented in [TS13].

We present a dual framework to analyze the proposed tensor completion formulation. This leads to novel fixed-rank formulations, for which we exploit the Riemannian framework to develop scalable trust region algorithms. In experiments, our algorithm TR-MM obtains better generalization performance and is more robust to outliers than state-of-the-art low-rank tensor completion algorithms. Overall both TR-MM and TR-LS algorithms achieve better performance on various completions tasks.

Acknowledgement

References

[ADKM11]   E. Acar, D. M. Dunlavy, T. G. Kolda, and M. Mørup, *Scalable tensor factorizations for incomplete data*, Chemometrics and Intelligent Laboratory Systems **106** (2011), no. 1, 41–56.

[AEP06]   A. Argyriou, T. Evgeniou, and M. Pontil, *Multi-task feature learning*, NIPS, 2006.

[AMS08]   P.-A. Absil, R. Mahony, and R. Sepulchre, *Optimization algorithms on matrix manifolds*, Princeton University Press, 2008.

[Ber99]   D. Bertsekas, *Nonlinear programming*, Athena Scientific, 1999.

[BMAS14]   N. Boumal, B. Mishra, P.-A. Absil, and R. Sepulchre, *Manopt, a Matlab toolbox for optimization on manifolds*, Journal of Machine Learning Research **15** (2014), no. Apr, 1455–1459.

[BS00]   J. F. Bonnans and A. Shapiro, *Perturbation analysis of optimization problems*, Springer-Verlag, 2000.

[BUGD$^+$13] A. Bordes, N. Usunier, A. Garcia-Duran, J. Weston, and O. Yakhnenko, *Translating embeddings for modeling multi-relational data*, NIPS, 2013.

[CC14]   C. F. Caiafa and A. Cichocki, *Stable, robust, and super fast reconstruction of tensors using multi-way projections*, IEEE Transactions on Signal Processing **63** (2014), no. 3, 780–793.

[CMR09]   C. Cortes, M. Mohri, and A. Rostamizadeh, *L2 regularization for learning kernels*, UAI, 2009.

[CPZ$^+$17a] A. Cichocki, H.-A. Phan, Q. Zhao, N. Lee, I. Oseledets, and D. P. Mandic, *Tensor networks for dimensionality reduction and large-scale optimization: Part 1 low-rank tensor decompositions*, Foundations and Trends in Machine Learning **9** (2017), no. 4–5, 249–429.

[CPZ$^+$17b] _____, *Tensor networks for dimensionality reduction and large-scale optimization: Part 2 applications and future perspectives*, Foundations and Trends in Machine Learning **9** (2017), no. 6, 431–673.

[CYZ$^+$16] H. Cheng, Y. Yu, X. Zhang, E. Xing, and D. Schuurmans, *Scalable and sound low-rank tensor learning*, AISTATS, 2016.

[EAC15]   B. Ermiş, E. Acar, and A. T. Cemgil, *Link prediction in heterogeneous data via generalized coupled tensor factorization*, KDD, 2015.

[EAS98]   A. Edelman, T.A. Arias, and S.T. Smith, *The geometry of algorithms with orthogonality constraints*, SIAM Journal on Matrix Analysis and Applications **20** (1998), no. 2, 303–353.

[FJ15]   M. Filipović and A. Jukić, *Tucker factorization with missing data with application to low-n-rank tensor completion*, Multidimensional Systems and Signal Processing (2015).

[FNA04]   D. H. Foster, S. M. C. Nascimento, and K. Amano, *Information limits on neural identification of colored surfaces in natural scenes*, Visual Neuroscience **21** (2004), no. 3, 331–336, URL: https://personalpages.manchester.ac.uk/staff/d.h.foster/.

[GYK17]   X. Guo, Q. Yao, and J. T. Kwok, *Efficient sparse low-rank tensor completion using the frank-wolfe algorithm*, AAAI, 2017, URL: https://github.com/quanmingyao/FFWTensor.

[HK15]   F. M. Harper and J. A. Konstan, *The movielens datasets: History and context*, ACM Transactions on Interactive Intelligent Systems **5** (2015), no. 4, 19:1–19:19, URL: http://files.grouplens.org/datasets/movielens/ml-10m-README.html.

[JBAS10]   M. Journée, F. Bach, P.-A. Absil, and R. Sepulchre, *Low-rank optimization on the cone of positive semidefinite matrices*, SIAM Journal on Optimization **20** (2010), no. 5, 2327–2351.

[JM18]   P. Jawanpuria and B. Mishra, *A unified framework for structured low-rank matrix learning*, ICML, 2018.

[JVN14]   P. Jawanpuria, M. Varma, and J. S. Nath, *On p-norm path following in multiple kernel learning for non-linear feature selection*, ICML, 2014.

[KB09]   T. G. Kolda and B. W. Bader, *Tensor decompositions and applications*, SIAM Review **51** (2009), no. 3, 455–500.

[KM16]   H. Kasai and B. Mishra, *Low-rank tensor completion: a Riemannian manifold preconditioning approach*, ICML, 2016, URL: https://bamdevmishra.in/codes/tensorcompletion/.

[KSV14]   D. Kressner, M. Steinlechner, and B. Vandereycken, *Low-rank tensor completion by Riemannian optimization*, BIT Numerical Mathematics **54** (2014), no. 2, 447–468, URL: anchp.epfl.ch/geomCG.

[LAAW16]   X.-Y. Liu, S. Aeron, V. Aggarwal, and X. Wang, *Low-tubal-rank tensor completion using alternating minimization*, SPIE Conference on Defense and Security, 2016.

[LMWY13]   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 **35** (2013), no. 1, 208–220, URL: http://www.cs.rochester.edu/u/jliu/code/TensorCompletion.zip.

[LZ11]   L. Lü and T. Zhou, *Link prediction in complex networks: A survey*, Physica A: statistical mechanics and its applications **390** (2011), no. 6, 1150–1170.

[PTJY10]   T. K. Pong, P. Tseng, S. Ji, and J. Ye, *Trace norm regularization: Reformulations, algorithms, and multi-task learning*, SIAM Journal on Optimization **20** (2010), no. 6, 3465–3489.

[RpABbP13] B. Romera-paredes, H. Aung, N. Bianchi-berthouze, and M. Pontil, *Multilinear multi-task learning*, ICML, 2013.

[SDLS14]   M. Signoretto, Q. T. Dinh, L. D. Lathauwer, and J. A. K. Suykens, *Learning with tensors: a framework based on convex optimization and spectral regularization*, Machine Learning **94** (2014), no. 3, 303–351.

[SI13]   H. Sato and T. Iwai, *A new, globally convergent Riemannian conjugate gradient method*, Optimization: A Journal of Mathematical Programming and Operations Research **64** (2013), no. 4, 1011–1031.

[SKM17]   H. Sato, H. Kasai, and B. Mishra, *Riemannian stochastic variance reduced gradient*, Tech. report, arXiv preprint arXiv:1702.05594, 2017.

[SNM08]   P. Symeonidis, A. Nanopoulos, and Y. Manolopoulos, *Tag recommendations based on tensor dimensionality reduction*, RecSys, 2008.

[Suz11]   T. Suzuki, *Unifying framework for fast learning rate of non-sparse multiple kernel learning*, NIPS, 2011.

[TCP$^+$15]   K. Toutanova, D. Chen, P. Pantel, H. Poon, P. Choudhury, and M. Gamon, *Representing text for joint embedding of text and knowledge bases*, EMNLP, 2015, URL: http://kristinatoutanova.com/.

[THK10]   R. Tomioka, K. Hayashi, and H. Kashima, *Estimation of low-rank tensors via convex optimization*, Tech. report, arXiv preprint arXiv:1010.0789, 2010.

[TS13]   R. Tomioka and T. Suzuki, *Convex tensor decomposition via structured schatten norm regularization*, NIPS, 2013, URL: http://tomioka.dk/softwares/.

[TSHK11]   R. Tomioka, T. Suzuki, K. Hayashi, and H. Kashima, *Statistical performance of convex tensor decomposition*, NIPS, 2011.

[TWL09]   L. Tang, X. Wang, and H. Liu, *Uncovering groups via heterogeneous interaction analysis*, ICDM, 2009, URL: http://leitang.net/heterogeneous_network.html.

[WST14]   K. Wimalawarne, M. Sugiyama, and R. Tomioka, *Multitask learning meets tensor factorization: task imputation via convex optimization*, NIPS, 2014.

[XJ12]   Y. Xin and T. Jaakkola, *Primal-dual methods for sparse constrained matrix completion*, AISTATS, 2012.

[ZEA$^+$14]   Z. Zhang, G. Ely, S. Aeron, N. Hao, and M. E. Kilmer, *Novel methods for multilinear data completion and de-noising based on tensor-svd.*, CVPR, 2014, URL: http://www.ece.tufts.edu/ shuchin/software.html.

[ZRS16]   H. Zhang, S. J. Reddi, and S. Sra, *Riemannian svrg: Fast stochastic optimization on Riemannian manifolds*, NIPS, 2016.

[ZZC15]   Q. Zhao, L. Zhang, and A. Cichocki, *Bayesian CP factorization of incomplete tensors with automatic rank determination*, IEEE Transactions on Pattern Analysis and Machine Intelligence **37** (2015), no. 9, 1751–1763, URL: https://github.com/qbzhao/BCPF.

A Proof of Lemma 3

Proof. Let $R(\mathcal{W}) = (\sum_k \|\mathcal{W}_k^{(k)}\|_*^2)^{\frac{1}{2}}$, from the optimality of $\hat{\mathcal{W}}$, we have

$$\mathcal{Y} - \hat{\mathcal{W}} \in \frac{1}{\lambda} \partial R(\hat{\mathcal{W}})^2, \tag{24}$$

where $\partial R(\hat{\mathcal{W}})^2$ is the subdifferential of $R(\mathcal{W})^2$ at $\mathcal{W} = \hat{\mathcal{W}}$.
From the triangle inequality, we obtain

$$\left\| \hat{\mathcal{W}} - \mathcal{W}^* \right\|_F \leq \left\| \hat{\mathcal{W}} - \mathcal{Y} \right\|_F + \|\mathcal{E}\|_F.$$

First term in the right hand side of the above inequality satisfies

$$\left\| \hat{\mathcal{W}} - \mathcal{Y} \right\|_F \leq \sqrt{\min_k n_k} \left\| \hat{\mathcal{W}} - \mathcal{Y} \right\|_{S_{\infty/2}}, \tag{25}$$

this can be seen from the following inequalities. For any tensor $\mathcal{W}$,

$$\|\mathcal{W}\|_F^2 \leq \min_k n_k \|\mathcal{W}_k\|_{S_\infty}^2 \leq \min_k n_k \sum_{k=1}^K \|\mathcal{W}_k\|_{S_\infty}^2.$$

For any $\mathcal{X} \in \partial R(\mathcal{W})^2$, we have $R_*(\mathcal{X})^2 \leq 1$, where $R_*(\mathcal{X})^2 := \|\mathcal{X}\|_{S_{\infty/2}}^2$ is the dual of $R(\mathcal{X})^2$. Using this in (25) ,

$$\left\| \hat{\mathcal{W}} - \mathcal{Y} \right\|_F \leq \frac{1}{\lambda} \sqrt{\min_k n_k}.$$

Following similar analysis and from the assumption we have

$$\|\mathcal{E}\|_F \leq \sqrt{\min_k n_k} \|\mathcal{E}\|_{S_{\infty/2}} \leq \frac{1}{\lambda} \sqrt{\min_k n_k},$$

$\square$

B Optimization on spectrahedron

We are interested in the optimization problem of the form

$$\min_{\Theta \in \mathcal{P}^d} \quad f(\Theta), \tag{A1}$$

where $\mathcal{P}^d$ is the set of $d \times d$ positive semi-definite matrices with unit trace and $f : \mathcal{P}^d \to \mathbb{R}$ is a smooth function. A specific interest is when we seek matrices of rank r. Using the parameterization $\Theta = \mathbf{U}\mathbf{U}^\top$, the problem (A1) is formulated as

$$\min_{\mathbf{U} \in \mathcal{S}_r^d} \quad f(\mathbf{U}\mathbf{U}^\top), \tag{A2}$$

where $\mathcal{S}_r^d := \{\mathbf{U} \in \mathbb{R}^{d \times r} : \|\mathbf{U}\|_F = 1\}$, which is called the spectrahedron manifold [JBAS10]. It should be emphasized the objective function in (A2) is *invariant* to the post multiplication of $\mathbf{U}$ with orthogonal matrices of size $r \times r$, i.e., $\mathbf{U}\mathbf{U}^\top = \mathbf{U}\mathbf{Q}(\mathbf{U}\mathbf{Q})^\top$ for all $\mathbf{Q} \in \mathcal{O}(r)$, which is the set of orthogonal matrices of size $r \times r$ such that $\mathbf{Q}\mathbf{Q}^\top = \mathbf{Q}^\top\mathbf{Q} = \mathbf{I}$. An implication of the this

Table 8: Matrix characterization of notions on the quotient manifold $\mathcal{S}_r^d/\mathcal{O}(r)$.

Matrix representation of an element	$\mathbf{U}$
Total space $\mathcal{S}_r^d$	$\{\mathbf{U} \in \mathbb{R}^{d\times r} : \|\mathbf{U}\|_F = 1\}$
Group action	$\mathbf{U} \mapsto \mathbf{U}\mathbf{Q}$, where $\mathbf{Q} \in \mathcal{O}(r)$.
Quotient space $\mathcal{M}$	$\mathcal{S}_r^d/\mathcal{O}(r)$
Tangent vectors in the total space $\mathcal{S}_r^d$ at $\mathbf{U}$	$\{\mathbf{Z} \in \mathbb{R}^{d\times r} : \mathrm{trace}(\mathbf{Z}^\top\mathbf{U}) = 0\}$
Metric between the tangent vector $\xi_{\mathbf{U}}, \eta_{\mathbf{U}} \in T_{\mathbf{U}}\mathcal{S}_r^d$	$\mathrm{trace}(\xi_{\mathbf{U}}^\top \eta_{\mathbf{U}})$
Vertical tangent vectors at $\mathbf{U}$	$\{\mathbf{U}\boldsymbol{\Lambda} : \boldsymbol{\Lambda} \in \mathbb{R}^{r\times r}, \boldsymbol{\Lambda}^\top = -\boldsymbol{\Lambda}\}$
Horizontal tangent vectors	$\{\xi_{\mathbf{U}} \in T_{\mathbf{U}}\mathcal{S}_r^d : \xi_{\mathbf{U}}^\top\mathbf{U} = \mathbf{U}^\top\xi_{\mathbf{U}}\}$

observation is that the minimizers of (A2) are no longer isolated in the matrix space, but are isolated in the quotient space, which is the set of equivalence classes $[\mathbf{U}] \coloneqq \{\mathbf{U}\mathbf{Q} : \mathbf{Q}\mathbf{Q}^\top = \mathbf{Q}^\top\mathbf{Q} = \mathbf{I}\}$. Consequently, the search space is

$$\mathcal{M} \coloneqq \mathcal{S}_r^d/\mathcal{O}(r). \tag{A3}$$

In other words, the optimization problem (A2) has the structure of optimization on the *quotient* manifold, i.e.,

$$\min_{[\mathbf{U}]\in\mathcal{M}} \quad f([\mathbf{U}]), \tag{A4}$$

but numerically, by necessity, algorithms are implemented in the matrix space $\mathcal{S}_r^d$, which is also called the *total space*.

Below, we briefly discuss the manifold ingredients and their matrix characterizations for (A4). Specific details of the spectrahedron manifold are discussed in [JBAS10]. A general introduction to manifold optimization and numerical algorithms on manifolds are discussed in [AMS08].

B.1 Tangent vector representation as horizontal lifts

Since the manifold $\mathcal{M}$, defined in (A3), is an abstract space, the elements of its tangent space $T_{[\mathbf{U}]}\mathcal{M}$ at $[\mathbf{U}]$ also call for a matrix representation in the tangent space $T_{\mathbf{U}}\mathcal{S}_r^d$ that respects the equivalence relation $\mathbf{U}\mathbf{U}^\top = \mathbf{U}\mathbf{Q}(\mathbf{U}\mathbf{Q})^\top$ for all $\mathbf{Q} \in \mathcal{O}(r)$. Equivalently, the matrix representation of $T_{[\mathbf{U}]}\mathcal{M}$ should be restricted to the directions in the tangent space $T_{\mathbf{U}}\mathcal{S}_r^d$ on the total space $\mathcal{S}_r^d$ at $\mathbf{U}$ that do not induce a displacement along the equivalence class $[\mathbf{U}]$. In particular, we decompose $T_{\mathbf{U}}\mathcal{S}_r^d$ into complementary subspaces, the *vertical* $\mathcal{V}_{\mathbf{U}}$ and *horizontal* $\mathcal{H}_{\mathbf{U}}$ subspaces, such that $\mathcal{V}_{\mathbf{U}} \oplus \mathcal{H}_{\mathbf{U}} = T_{\mathbf{U}}\mathcal{S}_r^d$.

The vertical space $\mathcal{V}_{\mathbf{U}}$ is the tangent space of the equivalence class $[\mathbf{U}]$. On the other hand, the horizontal space $\mathcal{H}_{\mathbf{U}}$, which is any complementary subspace to $\mathcal{V}_{\mathbf{U}}$ in $T_{\mathbf{U}}\mathcal{S}_r^d$, provides a valid matrix representation of the abstract tangent space $T_{[\mathbf{U}]}\mathcal{M}$. An abstract tangent vector $\xi_{[\mathbf{U}]} \in T_{[\mathbf{U}]}\mathcal{M}$ at $[\mathbf{U}]$ has a unique element in the horizontal space $\xi_{\mathbf{U}} \in \mathcal{H}_{\mathbf{U}}$ that is called its *horizontal lift*. Our

specific choice of the horizontal space is the subspace of $T_{\mathbf{U}}\mathcal{S}_r^d$ that is the *orthogonal complement* of $\mathcal{V}_x$ in the sense of a *Riemannian metric*.

The Riemannian metric at a point on the manifold is a inner product that is defined in the tangent space. An additional requirement is that the inner product needs to be *invariant* along the equivalence classes [AMS08, Chapter 3]. One particular choice of the Riemannian metric on the total space $\mathcal{S}_r^d$ is

$$\langle \xi_{\mathbf{U}}, \eta_{\mathbf{U}} \rangle_{\mathbf{U}} := \text{trace}(\xi_{\mathbf{U}}^\top \eta_{\mathbf{U}}), \tag{A5}$$

where $\xi_{\mathbf{U}}, \eta_{\mathbf{U}} \in T_{\mathbf{U}}\mathcal{S}_r^d$. The choice of the metric (A5) leads to a natural choice of the metric on the quotient manifold, i.e.,

$$\langle \xi_{[\mathbf{U}]}, \eta_{[\mathbf{U}]} \rangle_{[\mathbf{U}]} := \text{trace}(\xi_{\mathbf{U}}^\top \eta_{\mathbf{U}}),$$

where $\xi_{[\mathbf{U}]}$ and $\eta_{[\mathbf{U}]}$ are abstract tangent vectors in $T_{[\mathbf{U}]}\mathcal{M}$ and $\xi_{\mathbf{U}}$ and $\eta_{\mathbf{U}}$ are their horizontal lifts in the total space $\mathcal{S}_r^d$, respectively. Endowed with this Riemannian metric, the quotient manifold $\mathcal{M}$ is called a *Riemannian* quotient manifold of $\mathcal{S}_r^d$.

Table 8 summarizes the concrete matrix operations involved in computing horizontal vectors.

Additionally, starting from an arbitrary matrix (an element in the ambient dimension $\mathbb{R}^{d \times r}$), two linear projections are needed: the first projection $\Psi_{\mathbf{U}}$ is onto the tangent space $T_{\mathbf{U}}\mathcal{S}_r^d$ of the total space, while the second projection $\Pi_{\mathbf{U}}$ is onto the horizontal subspace $\mathcal{H}_{\mathbf{U}}$.

Given a matrix $\mathbf{Z} \in \mathbb{R}^{d \times r}$, the projection operator $\Psi_{\mathbf{U}} : \mathbb{R}^{d \times r} \to T_{\mathbf{U}}\mathcal{S}_r^d : \mathbf{Z} \mapsto \Psi_{\mathbf{U}}(\mathbf{Z})$ on the tangent space is defined as

$$\Psi_{\mathbf{U}}(\mathbf{Z}) = \mathbf{Z} - \text{trace}(\mathbf{Z}^\top \mathbf{U})\mathbf{U}. \tag{A6}$$

Given a tangent vector $\xi_{\mathbf{U}} \in T_{\mathbf{U}}\mathcal{S}_r^d$, the projection operator $\Pi_{\mathbf{U}} : T_{\mathbf{U}}\mathcal{S}_r^d \to \mathcal{H}_{\mathbf{U}} : \xi_{\mathbf{U}} \mapsto \Pi_{\mathbf{U}}(\xi_{\mathbf{U}})$ on the horizontal space is defined as

$$\Pi_{\mathbf{U}}(\xi_{\mathbf{U}}) = \xi_{\mathbf{U}} - \mathbf{U}\mathbf{\Lambda}, \tag{A7}$$

where $\mathbf{\Lambda}$ is the solution to the *Lyapunov* equation

$$(\mathbf{U}^\top \mathbf{U})\mathbf{\Lambda} + \mathbf{\Lambda}(\mathbf{U}^\top \mathbf{U}) = \mathbf{U}^\top \xi_{\mathbf{U}} - \xi_{\mathbf{U}}^\top \mathbf{U}.$$

B.2 Retractions from horizontal space to manifold

An iterative optimization algorithm involves computing a search direction (*e.g.,* the gradient direction) and then moving in that direction. The default option on a Riemannian manifold is to move along geodesics, leading to the definition of the exponential map. Because the calculation of the exponential map can be computationally demanding, it is customary in the context of manifold optimization to relax the constraint of moving along geodesics. The exponential map is then relaxed to a *retraction* operation, which is any map $R_{\mathbf{U}} : \mathcal{H}_{\mathbf{U}} \to \mathcal{S}_r^d : \xi_{\mathbf{U}} \mapsto R_{\mathbf{U}}(\xi_{\mathbf{U}})$ that locally approximates the exponential map on the manifold [AMS08, Definition 4.1.1]. On the spectrahedron manifold, a natural retraction of choice is

$$R_{\mathbf{U}}(\xi_{\mathbf{U}}) := (\mathbf{U} + \xi_{\mathbf{U}})/\|\mathbf{U} + \xi_{\mathbf{U}}\|_F,$$

where $\| \cdot \|_F$ is the Frobenius norm and $\xi_{\mathbf{U}}$ is a search direction on the horizontal space $\mathcal{H}_{\mathbf{U}}$.

An update on the spectrahedron manifold is, thus, based on the update formula $\mathbf{U}_+ = R_{\mathbf{U}}(\xi_{\mathbf{U}})$.

B.3 Riemannian gradient and Hessian computations

The choice of the invariant metric (A5) and the horizontal space turns the quotient manifold $\mathcal{M}$ into a *Riemannian submersion* of $(\mathcal{S}_r^d, \langle \cdot, \cdot \rangle)$. As shown by [AMS08], this special construction allows for a convenient matrix characterization of the gradient and the Hessian of a function on the abstract manifold $\mathcal{M}$.

The matrix characterization of the Riemannian gradient is

$$\mathrm{grad}_{\mathbf{U}}\, f = \Psi_{\mathbf{U}}(\nabla_{\mathbf{U}} f), \tag{A8}$$

where $\nabla_{\mathbf{U}} f$ is the Euclidean gradient of the objective function f and $\Psi_{\mathbf{U}}$ is the tangent space projector (A6).

An iterative algorithm that exploits second-order information usually requires the Hessian applied along a search direction. This is captured by the Riemannian Hessian operator Hess, whose matrix characterization, given a search direction $\xi_{\mathbf{U}} \in \mathcal{H}_{\mathbf{U}}$, is

$$\mathrm{Hess}_{\mathbf{U}}[\xi_{\mathbf{U}}] = \Pi_{\mathbf{U}}\Big(\mathrm{D}\nabla f[\xi_{\mathbf{U}}] - \mathrm{trace}((\nabla_{\mathbf{U}} f)^{\top}\mathbf{U})\xi_{\mathbf{U}} - \mathrm{trace}((\nabla_{\mathbf{U}} f)^{\top}\xi_{\mathbf{U}} + (\mathrm{D}\nabla f[\xi_{\mathbf{U}}])^{\top}\mathbf{U})\mathbf{U}\Big), \tag{A9}$$

where $\mathrm{D}\nabla f[\xi_{\mathbf{U}}]$ is the directional derivative of the Euclidean gradient $\nabla_{\mathbf{U}} f$ along $\xi_{\mathbf{u}}$ and $\Pi_{\mathbf{U}}$ is the horizontal space projector (A7).

Finally, the formulas in (A8) and (A9) that the Riemannian gradient and Hessian operations require only the expressions of the standard (Euclidean) gradient of the objective function f and the directional derivative of this gradient (along a given search direction) to be supplied.