Tensor Completion via Integer Optimization

Xin Chen [* 1] **Sukanya Kudva** [* 1] **Yongzheng Dai** [2] **Anil Aswani** [1] **Chen Chen** [2]

Abstract

The main challenge with the tensor completion problem is a fundamental tension between computation power and the information-theoretic sample complexity rate. Past approaches either achieve the information-theoretic rate but lack practical algorithms to compute the corresponding solution, or have polynomial-time algorithms that require an exponentially-larger number of samples for low estimation error. This paper develops a novel tensor completion algorithm that resolves this tension by achieving both provable convergence (in numerical tolerance) in a linear number of oracle steps and the information-theoretic rate. Our approach formulates tensor completion as a convex optimization problem constrained using a gauge-based tensor norm, which is defined in a way that allows the use of integer linear optimization to solve linear separation problems over the unit-ball in this new norm. Adaptations based on this insight are incorporated into a Frank-Wolfe variant to build our algorithm. We show our algorithm scales-well using numerical experiments on tensors with up to ten million entries.

1. Introduction

A tensor is a multilinear operator, and it can be represented as an array of numbers referenced by multiple indices. Though vectors and matrices are special cases of tensors, tensors with three or more indices pose unique challenges. Many tensor problems are NP-hard (Hillar & Lim, 2013), such as computing: rank, singular values, and nuclear norm.

In this paper, we address one such difficult problem called *tensor completion*. Here, a small subset of tensor entries are observed – possibly with noise. Under an assumption of low-rankness, the problem is to fill-in the remaining, un-observed entries – and remove noise, if any. Since modern datasets are often multidimensional, there are many applications of tensor completion (Song et al., 2019), including: recommendation systems (Ge et al., 2016; Karatzoglou et al., 2010), information diffusion (Zafarani et al., 2014), regression (Aswani, 2016), computer vision (Duarte & Baraniuk, 2012; Signoretto et al., 2011; Tan et al., 2013), and bioinformatics (Bazerque et al., 2013; Acar et al., 2011).

1.1. Related Work

The information-theoretic rate for estimation error is $\sqrt{k \cdot \sum_i r_i / n}$ for a tensor completion problem, where: k is tensor rank, r_i is the i-th dimension of the tensor, and n is the number of samples (Gandy et al., 2011). Previous works have tried to characterize the tradeoff between complexity and the information-theoretic rate. For example, Barak & Moitra (2016) uses the Rademacher complexity of a sum-of-squares hierarchy to suggest a suitable norm for tensor completion, and observe a gap between what can be achieved information theoretically and what is attained by their computationally efficient method.

Initial work on tensor completion used decomposition methods based on CP and Tucker decomposition (Tomasi & Bro, 2005). Other approaches accounted for robustness to outliers and data corruptions (Javed et al., 2015; Jain et al., 2017), or imposed fixed-rank constraints (Kressner et al., 2014). Alternative approaches used various tensor norms as a convex surrogate for tensor rank (Montanari & Sun, 2018; Yuan & Zhang, 2016; Barak & Moitra, 2016). Our approach falls into this category, but here we define a novel norm.

Past approaches either traded information-theoretic rate for a computationally efficient algorithm (Montanari & Sun, 2018; Barak & Moitra, 2016); likewise, heuristics have been developed to compute non-certifiably-optimal solutions to such NP-hard problems whose optimal solutions (if found) in principal achieve the information-theoretic rate (Kressner et al., 2014; Yuan & Zhang, 2016; Jain et al., 2017). In contrast, we present a globally convergent algorithm that attains the information-theoretic rate (hence data efficiency), while certifying optimal solutions on instances on general tensors of sizes up to $10^{\times 7}$ within minutes.

Certain algorithms for special cases of tensor completion have achieved the information-theoretic rate through practi-

---

[*]Equal contribution [1]Industrial Engineering and Operations Research, University of California Berkeley, Berkeley, California, USA [2]Integrated Systems Engineering, The Ohio State University, Columbus, Ohio, USA. Correspondence to: Sukanya Kudva <sukanya_kudva@berkeley.edu>, Xin (Jennifer) Chen <xin.chen2024@berkeley.edu>.

cal computation. Aswani (2016) studied tensor completion for rank-1 non-negative tensors, formulating it as conic optimization with exponential constraints. Symmetric tensor completion was also studied, where Rao et al. (2015) used a variant of the Frank-Wolfe algorithm and Cai et al. (2019) used two-stage non-convex optimization. In our previous work (Bugg et al., 2022), we proposed an algorithm based on first-order optimization, that uses integer optimization for a weak separation oracle, for the special case of nonnegative tensors. We also explored accelerated versions of this integer-optimization-based algorithm for nonnegative tensors in (Pan et al., 2023).

1.2. Contribution

In this paper, we design an algorithm for tensor completion of general tensors, which converges to a global optimum with a linear number of oracle calls and satisfies the information-theoretic rate. In past literature, tensor completion algorithms that perform both information-theoretically and computationally-well have been achieved only in special cases (e.g., nonnegative tensors, symmetric orthogonal tensors).

Our algorithm is a step in this direction for the case of general tensors. Building upon applications of various special-case tensor completion algorithms (Li et al., 2008; Aswani, 2016; Liu et al., 2012; Zhang et al., 2019), our algorithm for tensor completion opens up new applications that require positive *and* negative entries. Examples include social computing (Song et al., 2019), the moment method on multivariate distribution (Montanari & Sun, 2018), and healthcare applications (Gandy et al., 2011; Bazerque et al., 2013).

The main idea behind our algorithm is to define the tensor completion problem using a gauge norm, and then use a Frank-Wolfe-like first-order optimization algorithm to solve the newly defined convex optimization formulation. We define the gauge norm by constructing a convex polytope with its vertices as rank-1 tensors. The rank-1 tensor vertices help define the convex polytope using integer linear constraints. We relate the gauge norm to tensor rank, and analyze the norm's computational and statistical complexity to NP-hard and low Rademacher complexity respectively.

Consequently, the tensor completion problem using gauge norm is also NP-hard to solve to arbitrary accuracy. Nevertheless, since the formulation is a convex optimization, we design an algorithm using Blended Conditional Gradients (BCG) (Braun et al., 2019). We construct a weak separation oracle for BCG using an integer linear optimization formulation and an additional heuristic to accelerate computation. Our numerical experiments demonstrate that this algorithm achieves the information-theoretic rate and is efficient for tensors with as large as ten million entries.

Our paper is organized as follows. In Section 2, we go over basic notations and the tensor completion problem. Section 3 defines the gauge norm and establishes its relationship with tensor rank. We further analyze the norm complexity in Section 4. Then in Section 5, we discuss the complexity of the tensor completion problem with gauge norm, and design an algorithm. Finally, Section 6 summarizes our numerical experiments, followed by the conclusion in Section 7.

1.3. Going from Nonnegative to General Tensors

Since the approach in our previous work (Bugg et al., 2022) is closely related to this paper, here we provide a brief discussion about why going from nonnegative to general tensors requires the design and analysis of a new tensor completion algorithm:

The key idea in our previous work was to define a gauge norm using a 0-1 polytope that represents the convex hull of all rank-1 nonnegative tensors whose maximum entry is 1. The polytope was designed in such a way that linear separation problems over the polytope could be written using integer linear constraints. The approach can be illustrated by an example: For the set $\{\zeta = \prod_{k=1}^{p} \nu_k : \nu_k \in \{0, 1\}\}$ defined using a product, its linearization is given by $\{\zeta : 0 \le \zeta \le \nu_k, \sum_k \nu_k + (1 - p) \le \zeta, \nu_k \in \{0, 1\}\}$. A direct linearization is possible because of the properties that: $\zeta = 0$ if any single $\nu_k = 0$, and $\zeta = 1$ if all $\nu_k = 1$.

The natural generalization of this idea, which we pursue in this paper, is to design a gauge norm using a polytope whose vertices have ± 1 entries and that represents the convex hull of all (general) rank-1 tensors whose maximum/minimum entry is ± 1. However, the above property used for linearization no longer holds when considering the new polytope. This can be shown via an example: For the set $\{\zeta = \prod_{k=1}^{p} \nu_k : \nu_k \in \{-1, +1\}\}$, a direct linearization is not possible because $\zeta = +1$ if an even number of $\nu_k = -1$, and $\zeta = -1$ if an odd number of $\nu_k = -1$. The generalization is fundamentally different because it involves the parity (i.e., odd or even) of the underlying items being multiplied. This is much more challenging and requires new computational design and theoretical analysis.

General tensors also have well-posedness issues that do not occur for nonnegative tensors. In particular, the best low-rank approximation problem is not well-posed for tensors (de Silva & Lim, 2008); this is related to a well-known example that a sequence of rank-2 tensors can be constructed whose limit has rank-3. In contrast, the best low-rank approximation problem *is* well-posed for nonnegative tensors (Qi et al., 2016). These phenomena create a technical challenge in determining whether a new gauge norm as defined above can act as a convex surrogate for the rank of tensors.

2. Preliminaries

For a tensor $\psi \in \mathbb{R}^{r_1 \times \cdots \times r_p}$ of order p, we use the notation $\psi_x := \psi_{x_1,\ldots,x_p}$ to refer to an entry corresponding to the indices $x = (x_1, \ldots, x_p)$, where $x_i \in [r_i]$ and $[s] := \{1, \ldots, s\}$. Furthermore, we define some parameters that depend upon the tensor dimensions and order: $\rho = \sum_i r_i$, $\pi = \prod_i r_i$, and $\mathcal{R} = [r_1] \times \cdots \times [r_p]$.

By definition, a rank-1 tensor can be written as the tensor product of vectors, that is $\psi = \bigotimes_{k=1}^p \theta^k$, where $\theta^{(k)} \in \mathbb{R}^{r_k}$. Equivalently, each tensor entry is $\psi_x = \prod_{k=1}^p \theta_{x_k}^{(k)}$, where $\theta_{x_k}^{(k)}$ is the x_k-th element of vector $\theta^{(k)}$ for any index value $x_k \in [r_k]$. When obvious, we will use θ_{x_k} instead of $\theta_{x_k}^{(k)}$.

Let $\mathcal{B}_\lambda$ be the set of rank-1 tensors such that each entry of the tensor has absolute values less than or equal to $\lambda \in \mathbb{R}_+$:

$$\mathcal{B}_\lambda = \{\psi : \psi_x = \lambda \cdot \textstyle\prod_{k=1}^p \theta_{x_k}, \ \theta_{x_k} \in [-1,1] \text{ for } x \in \mathcal{R}\}, \tag{1}$$

Then, the rank of a tensor is defined to be the minimum number of rank-1 tensors required to represent it. Formally,

$$\mathrm{rank}(\psi) = \min\{q \mid \psi = \textstyle\sum_{k=1}^q \psi^k, \psi^k \in \mathcal{B}_\infty \text{ for } k \in [q]\}.$$

Using this definition of tensor rank, a CP decomposition of the tensor is given by $\psi = \sum_{k=1}^{\mathrm{rank}(\psi)} \psi^k$.

The tensor completion problem we consider begins with n observations of the tensor, which are denoted by the pairs $(x\langle i \rangle, y\langle i \rangle) \in \mathcal{R} \times \mathbb{R}$ for $i \in [n]$. Here, $y\langle i \rangle$ is the (possibly noisy) observation of the tensor entry $\psi_{x\langle i \rangle}$. We note that the $x\langle i \rangle$ are assumed to be independent and identically distributed in our model, which means that any given entry of the tensor may be observed multiple times within the n observations. Our approach is to solve the tensor completion problem using a least squares formulation:

$$\widehat{\psi} \in \arg\min_\psi \ \frac{1}{n} \sum_{i=1}^n \left(y\langle i \rangle - \psi_{x\langle i \rangle} \right)^2$$
$$\text{s.t. } \|\psi\|_\pm \leq \lambda \tag{2}$$

where the constraint uses a new norm $\|\psi\|_\pm$ that we design below. We show that our norm does in fact act as a convex surrogate for tensor rank, analyze its complexity, and demonstrate that it leads to an efficient computation algorithm for the tensor completion problem.

3. Gauge Norm for Tensors

We construct a norm for general tensors using a gauge function (Chandrasekaran et al., 2012; Jaggi, 2013; Bugg et al., 2022). To start with, consider rank-1 tensors. Let $\mathcal{S}_\lambda$ be the set of rank-1 tensors such that each entry of the tensor has an absolute value of some $\lambda \in \mathbb{R}_+$:

$$\mathcal{S}_\lambda = \{\psi : \psi_x = \lambda \cdot \textstyle\prod_{k=1}^p \theta_{x_k}, \ \theta_{x_k} \in \{-1,1\} \text{ for } x \in \mathcal{R}\}. \tag{3}$$

Our first step in defining a norm is to relate the convex hull of these sets of rank-1 tensors.

Proposition 3.1. *The convex hulls of these sets are the same, meaning we have* $\mathcal{C}_\lambda := \mathrm{conv}(\mathcal{B}_\lambda) = \mathrm{conv}(\mathcal{S}_\lambda)$.

Proof. The proof for this proposition is similar to our proof of Proposition 2.1 in (Bugg et al., 2022), but where a multi-linear optimization problem is formulated by restricting the entries to $[-1, 1]$ instead of $[0, 1]$. $\square$

Remark 3.2. Note $\mathcal{B}_\lambda = \lambda \mathcal{B}_1$, $\mathcal{S}_\lambda = \lambda \mathcal{S}_1$, and $\mathcal{C}_\lambda = \lambda \mathcal{C}_1$.

3.1. Defining the Norm

$\mathcal{C}_\lambda$ is useful because it is a convex polytope with its vertices as the points in $\mathcal{S}_\lambda$. Our next step is to use $\mathcal{C}_\lambda$ to define a function that we prove in the next proposition is a gauge norm. The proof reveals the non-intuitive property that $\mathcal{C}_1$ has a non-empty interior and hence can be magnified (using $\lambda \mathcal{C}_1$) to cover all tensors. This property ensures that the below function is in fact a norm defined for all tensors.

Proposition 3.3. *The function defined as*

$$\|\psi\|_\pm := \inf\{\lambda \geq 0 \mid \psi \in \lambda \mathcal{C}_1\} \tag{4}$$

is a norm for all tensors $\psi \in \mathbb{R}^{r_1 \times \cdots \times r_p}$.

Proof. We use the result from Example 3.50 of (Rockafellar & Wets, 2009) to conclude that $\| \cdot \|_\pm$ is a norm. To apply the result, we check that the required conditions on $\mathcal{C}_1$ hold.

By definition $\mathcal{C}_1$ is convex, closed, and bounded. $\mathcal{C}_1$ is also symmetric since for every $a \in \mathcal{C}_1$, $-a \in \mathcal{C}_1$ is also true. To see this, let $a = \sum_i \lambda_i \psi_i$ where $\psi_i \in \mathcal{S}_1$, $\lambda_i \in [0, 1]$ and $\sum_i \lambda_i = 1$ and notice $-\psi_i \in \mathcal{S}_1$. Symmetry and convexity also ensure $0 \in \mathcal{C}_1$ since for any $a \in \mathcal{C}_1$, $\frac{1}{2}(a + (-a)) = 0$.

The final, non-trivial condition required is that $\mathcal{C}_1$ has a non-empty interior. To prove this, we use Theorem 2.4 of (Rockafellar, 2015) that the dimension of $\mathcal{C}_1$ is the maximum of dimensions of simplices included in it. We construct a simplex in $\mathcal{C}_1$ that has dimension $\pi = \Pi_i r_i$, and hence $\mathcal{C}_1$ has a dimension of at least π. But the flattened vector space of tensors also has dimension π; consequently, the dimension of $\mathcal{C}_1$ is at most π. Hence, $\mathcal{C}_1$ has a dimension π, which is the full dimension of the space, implying that the set $\mathcal{C}_1$ must have a non-empty interior.

The rest of the proof constructs such a simplex of dimension π. Consider a polytope $D = \mathrm{conv}(0 \cup \{d^x\}_{x \in \mathcal{R}})$. Here, for any $x = (x_1, \cdots, x_p) \in \mathcal{R}$, consider $d^x = \bigotimes_{k=1}^p \beta^{x_k}$ with each vector $\beta^{x_k} \in \mathbb{R}^{r_k}$ defined as

$$\beta^{x_k} = \begin{cases} \mathbb{1} & \text{if } x_k = 1, \\ \mathbb{f}_{x_k} & \text{if } x_k \neq 1, \end{cases}$$

3

where $\mathbb{1}$ is a vector of one's and $\mathbb{f}_j$ is a vector with -1 in position j and one's elsewhere. One can verify that $\{\beta^{x_k}\}_{x_k \in [r_k]}$ are linearly independent vectors and make a complete basis for $\mathbb{R}^{r_k}$. Since the tensor product of linearly independent vectors gives linearly independent tensors, the tensors $\{d^x\}_{x \in \mathcal{R}}$ are all linearly independent. Consequently, $\{d^x - 0\}_{x \in \mathcal{R}}$ are linearly independent and $\{d^x\}_{x \in \mathcal{R}} \cup 0$ are affinely independent. By definition, the polytope D, which is a convex hull of $|\mathcal{R}| + 1 = \pi + 1$ affinely independent points, is a simplex of dimension π. Note that all the points are in $\mathcal{C}_1$, and hence so is the simplex.

With this, $\mathcal{C}_1$ is shown to satisfy all the required conditions, which means that the proposition holds. $\qquad\square$

Towards our end goal of developing a practical algorithm for tensor completion using this gauge norm, our next result provides an alternative characterization of the vertices of $\mathcal{C}_\lambda$. This result is important because it shows that these vertices can be represented by linear inequality constraints using (binary) integer variables.

Proposition 3.4. *Consider the set defined as*

$$
\begin{aligned}
\widehat{\mathcal{S}}_\lambda = \big\{\psi : \; &\psi_x = \lambda \cdot y_{x,1} & x \in \mathcal{R} \\
&y_{x,k} \geq (-\theta_{x_k} - y_{x,k+1} - 1) & k \in [p-1], x \in \mathcal{R} \\
&y_{x,k} \geq (\theta_{x_k} + y_{x,k+1} - 1) & k \in [p-1], x \in \mathcal{R} \\
&y_{x,k} \leq (\theta_{x_k} - y_{x,k+1} + 1) & k \in [p-1], x \in \mathcal{R} \\
&y_{x,k} \leq (-\theta_{x_k} + y_{x,k+1} + 1) & k \in [p-1], x \in \mathcal{R} \\
&y_{x,p} = \theta_{x_p} & x \in \mathcal{R} \\
&\theta_{x_k} \in \{-1, 1\} & x \in \mathcal{R} \\
&\theta^k \in \mathbb{R}^{r_k}, \; y_{x,k} \in \mathbb{R} & k \in [p], x \in \mathcal{R} \big\}.
\end{aligned}
$$

We have that $\widehat{\mathcal{S}}_\lambda = \mathcal{S}_\lambda$.

Proof. We show that the constraints defining the sets $\mathcal{S}_\lambda$ and $\widehat{\mathcal{S}}_\lambda$ are equivalent. Consider some $x \in \mathcal{R}$. From the definition of $\mathcal{S}_\lambda$ in (3), we have $\psi_x = \lambda \prod_{i=1}^{p} \theta_{x_k}$. Define $y_{x,k} = \prod_{i=k}^{p} \theta_{x_k}$, for $k \in [p]$ so that $\psi_x = \lambda y_{x,1}$. Or equivalently, in a recursive relationship, $y_{x,k} = \theta_{x_k} y_{x,k+1}$ for $k \in [p-1]$ and $y_{x,p} = \theta_{x_p}$. The recursive constraints can be thought of as a negated-XOR relation, and linearized by transformations for conjunctive and disjunctive statements (see Section 2.5 of (Conforti et al., 2014)). These linearized constraints correspond to constraints 2-5 in the definition of $\widehat{\mathcal{S}}_\lambda$ as given above. This shows that for each $x \in \mathcal{R}$, the tensor ψ_x is defined the same in both $\mathcal{S}_\lambda$ and $\widehat{\mathcal{S}}_\lambda$. $\qquad\square$

3.2. Relation between Gauge Norm and Tensor Rank

We next establish a relationship between tensor rank and our gauge norm, and we use this to argue that our norm is a meaningful constraint in the tensor completion problem.

The underlying issue is related to the fact that the best low-rank approximation problem is not well-posed for tensors (de Silva & Lim, 2008), which is in sharp contrast to the case of nonnegative tensors for which the best low-rank approximation problem is well-posed (Qi et al., 2016).

For the results in this subsection, we impose a regularity condition to eliminate such pathological behavior of tensors.

Assumption 3.5 (Regularity Condition). *Consider a class of tensors defined by the set*

$$
\begin{aligned}
\Gamma = \big\{\psi : \; &\exists \; \text{CP decomposition of } \psi \text{ with terms } \psi^k \text{ s.t.} \\
&\|\psi^k\|_{\max} \leq \|\psi\|_{\max} \text{ for } k \in [\text{rank}(\psi)]\big\}. \quad (5)
\end{aligned}
$$

This class is such that each tensor ψ has its largest entry at least as large as the largest entry of each CP term ψ^k.

Remark 3.6. A CP decomposition always exists for a finite-valued tensor, but it may not be unique. The class defined above asks that the regularity condition holds for at least one CP decomposition, but does not make any statement about holding for all the possible CP decompositions.

The following proposition suggests that the norm $\|\psi\|_\pm$, which is convex, can be a useful alternative to tensor rank.

Proposition 3.7. *For any $\psi \in \Gamma$ that satisfies Assumption 3.5, we have $\|\psi\|_{\max} \leq \|\psi\|_\pm \leq \text{rank}(\psi) \cdot \|\psi\|_{\max}$.*

Proof. The proof is similar to Proposition 2.4 of (Bugg et al., 2022). However, the right-side inequality requires the regularity assumption. Using the CP decomposition and the triangle inequality for norms,

$$
\|\psi\|_\pm \leq \sum_{k=1}^{\text{rank}(\psi)} \|\psi^k\|_\pm = \sum_{k=1}^{\text{rank}(\psi)} \|\psi^k\|_{\max},
$$
$$
\text{where } \psi^k \in \mathcal{B}_\infty
$$

The last equality follows by noting that $\|\psi^k\|_\pm = \|\psi^k\|_{\max}$ when $\psi^k \in \mathcal{B}_\infty$. Using $\|\psi^k\|_{\max} \leq \|\psi\|_{\max}$ from Assumption 3.5 gives the desired right-side inequality. $\qquad\square$

4. Complexity Analysis of Norm

We show that calculating the norm $\|\cdot\|_\pm$ is NP-hard. Despite this, it is still useful for tensor completion because it is defined using a convex polytope $\mathcal{C}_1$ whose vertices admit a convenient representation, as described in Proposition 3.4, that we will use to develop a practical algorithm.

Proposition 4.1 (**Computational complexity**). *The norm $\|\cdot\|_\pm$ is NP-hard to approximate to arbitrary accuracy.*

Proof. Note that $\|\varphi\|_\circ = \sup\{|\langle\varphi, \psi\rangle| \mid \|\psi\|_\pm \leq 1\} = \sup\{\langle\varphi, \psi\rangle \mid \psi \in \mathcal{C}_1\}$ is the dual norm for $\|\cdot\|_\pm$. The approximation of $\|\cdot\|_\circ$ can be reduced to approximation of the norm $\|\cdot\|_\pm$ in polynomial time from theorems 3

and 10 of (Friedland & Lim, 2016). Our main idea is to give a polynomial-time reduction of an NP-Hard problem to an approximation of $\| \cdot \|_\circ$. In particular, we prove that calculating the $\infty, 1$ subordinate matrix norm is polynomial-time reducible to $\sup\{\langle \varphi, \psi \rangle \mid \psi \in \mathcal{S}_1\}$. Without loss of generality, assume $p = 2$ and $d := r_1 = r_2$.

The decision version of $\sup\{\langle \varphi, \psi \rangle \mid \psi \in \mathcal{S}_1\}$ is:

Question: Does there exist $\theta_{x_k} \in \{-1, 1\}$ for all $x_k \in [d]$ with $k = 1, 2$ such that for a given L we have

$$\sum_{x_1=1}^d \sum_{x_2=1}^d \varphi_{x_1 x_2} \theta_{x_1} \theta_{x_2} \geq L \ ?$$

From Proposition 1 of (Rohn, 2000), we have $\|W\|_{\infty,1} = \{\max \sum_{i,j} W_{ij} x_i y_j \mid x_i, y_j \in \{-1, 1\}\}$ for a matrix W. The decision version of the $\infty, 1$ subordinate matrix norm for the special case of M-matrices can be written as:

Question: Does there exist $x_i, y_i \in \{-1, 1\}$ for all $i \in [d]$ such that for a given $L' \geq 0$ and a *symmetric, positive definite* matrix $W \in \mathbb{R}^{d \times d}$ satisfying $w_{ij} \leq 0$ for all $i \neq j$ and

$$\sum_{i=1}^d \sum_{j=1}^d w_{ij} x_i y_j \geq L' \ ?$$

Clearly, setting $L = L'$ and $\varphi = W$ is a valid polynomial time reduction. Since $\mathcal{C}_1 = \mathrm{conv}(\mathcal{S}_1)$ and $\langle \varphi, \psi \rangle$ is linear, $\|\varphi\|_\circ = \sup\{\langle \varphi, \psi \rangle \mid \psi \in \mathcal{S}_1\}$. The result now follows since approximately solving any ∞, p subordinate matrix norm, where $p \in [1, \infty)$, to arbitrary accuracy is NP-hard (Hendrickx & Olshevsky, 2010). In particular, Theorem 5 of (Rohn, 2000) shows that it is an NP-hard problem for M-matrices. $\qquad \square$

The same proof of Corollary 3.2 in (Bugg et al., 2022) combined with the above result establishes NP-completeness:

Corollary 4.2 (Bugg et al., 2022). *Given $K \in \mathbb{R}_+$ and $\psi \in \mathbb{R}^{r_1 \times \cdots \times r_p}$, it is NP-complete to determine if $\|\psi\|_\pm \leq K$.*

Rademacher complexity, from computational learning theory, is used to characterize the richness of a class of functions (Bartlett & Mendelson, 2002; Srebro et al., 2010). Roughly speaking, function classes with lower Rademacher complexity can be learned using less samples. From the following proposition, one can check that the norm $\| \cdot \|_\pm$ has an exponentially smaller Rademacher complexity than the max and Frobenius norms for tensors.

Proposition 4.3 (Stochastic Complexity). *We have $\mathsf{R}(\mathcal{C}_\lambda) \leq \mathsf{W}(\mathcal{C}_\lambda) \leq 2\lambda\sqrt{\rho/n}$, where $\mathsf{R}(\cdot)$ and $\mathsf{W}(\cdot)$ are the Rademacher and worst case Rademacher complexities.*

Proof. Rademacher complexity is computed using symmetric random variables $\sigma \in \{-1, +1\}$. It is for this reason that the class of functions $\mathcal{H}$, $-\mathcal{H}$, and $\mathcal{H} \cup -\mathcal{H}$ have the same Rademacher complexity. With this observation, the proof follows that of Proposition 3.3 of (Bugg et al., 2022). $\quad \square$

5. Algorithm for Tensor Completion

We now turn our attention towards numerical solution of the tensor completion problem (2) using our norm $\|\psi\|_\pm$.

5.1. Complexity Analysis of Tensor Completion

By interpreting the tensor completion problem (2) as a *convex aggregation* problem (Nemirovski, 2000; Tsybakov, 2003; Lecué, 2013) for a finite set of functions, one can arrive at the generalization bound for the solution. Interestingly, we had got the same generalization bounds for the special case of nonnegative tensors (Bugg et al., 2022). We believe this is because the proof for nonnegative tensors did not exploit the non-negativity, and hence could have led to non-tight generalization bounds. For completeness, we state these statistical guarantees in the following two results:

Proposition 5.1 (Lecué, 2013). *Suppose $|y| \leq b$ almost surely. Given any $\delta > 0$, with probability at least $1 - 4\delta$ we have that*

$$\mathbb{E}\big((y - \psi_x)^2\big) \leq \min_{\varphi \in \mathcal{C}_\lambda} \mathbb{E}\big((y - \varphi_x)^2\big) +$$
$$c_0 \cdot \max\big[b^2, \lambda^2\big] \cdot \max\big[\zeta_n, \tfrac{\log(1/\delta)}{n}\big], \quad (6)$$

where c_0 is an absolute constant and

$$\zeta_n = \begin{cases} \frac{2^\rho}{n}, & \text{if } 2^\rho \leq \sqrt{n} \\ \sqrt{\frac{1}{n} \log\left(\frac{e 2^\rho}{\sqrt{n}}\right)}, & \text{if } 2^\rho > \sqrt{n} \end{cases} \quad (7)$$

Corollary 5.2 (Bugg et al., 2022). *Suppose $\psi \in \Gamma$ is a tensor (satisfying Assumption 3.5) with $\mathrm{rank}(\psi) = k$ and $\|\psi\|_{\max} \leq \mu$. Under an additive noise model, if $(x\langle i \rangle, y\langle i \rangle)$ are independent and identically distributed with $|y\langle i \rangle - \varphi_{x\langle i \rangle}| \leq e$ almost surely and $\mathbb{E}y\langle i \rangle = \psi_{x\langle i \rangle}$. Then given any $\delta > 0$, with probability at least $1 - 4\delta$ we have*

$$\mathbb{E}\big((y - \widehat{\psi}_x)^2\big) \leq e^2 + c_0 \cdot (\mu k + e)^2 \cdot \max\big[\zeta_n, \tfrac{\log(1/\delta)}{n}\big], \quad (8)$$

where ζ_n is as in (7) and c_0 is an absolute constant.

Remark 5.3. The above result achieves the information-theoretic rate when the rank $k = O(1)$.

Unsurprisingly, the tensor completion problem (2) is NP-hard, because approximating the norm $\| \cdot \|_\pm$ is NP-hard and we had shown a polynomial-time reduction of the problem (2) to the NP-hard weak membership problem in our previous work (Bugg et al., 2022).

Proposition 5.4. *The tensor completion problem (2) is NP-hard to solve to arbitrary accuracy. Also, its decision version is NP-complete.*

Proof. The proof for this proposition is similar to the proof of Proposition 4.4 of (Bugg et al., 2022). $\qquad \square$

5.2. Numerical Algorithm

The problem (2) is a convex optimization formulation despite being NP-hard. This property enables the application of various first-order convex optimization algorithms. In particular, we use a variant of the Frank-Wolfe algorithm called *Blended Conditional Gradients* (BCG) (Braun et al., 2019). The main alternative approaches for iterative optimization seem to run into issues inherent to the implicit description of the feasible region; for instance, we do not have an effective barrier function available for interior point methods, nor do we have an efficient projection oracle that can be leveraged in projected gradient descent.

We find the BCG algorithm to be practically efficient for our problem since it requires a weak linear separation oracle leading to an early termination. When we design our weak separation oracle calls to an integer optimization problem, we explicitly define a tolerance for early termination. This works out well as integer optimization solvers usually discover near-optimal solutions fast and subsequently work hard to certify them. Further, the linear convergence of BCG is guaranteed when the feasible set is a polytope and the objective function is strictly convex. We construct a strictly convex problem by projecting the feasible space onto the set of unique, observed tensor entries like we did in (Bugg et al., 2022). Since the conditions for linear convergence are satisfied by our problem, we get an algorithm that terminates in a linear number of oracle steps.

The BCG algorithm makes subsequent iterate updates based on the gradient of the objective function at the current iterate. It needs a weak separation oracle to find a new vertex that reduces a gradient-based linear objective. The weak separation oracle has the same requirements as in Algorithm 1 in (Bugg et al., 2022). The output of the oracle should either give a vertex that accomplishes separation, or a certificate that separation is not possible. We design our weak separation oracle using two algorithms, an integer optimization problem in (9) and an alternating maximization heuristic.

Since integer optimization is likely to be more computationally expensive than the alternating heuristic, we first try the latter a few times with different initializations. We extend our heuristic in Algorithm 2 of (Bugg et al., 2022) to general tensors, and explore solutions by toggling between $\{-1, 1\}$ instead of $\{0, 1\}$. It exploits the fact that the objective is multi-linear and minimizes the objective in different dimensions, going one by one. It runs in polynomial time and has been seen to speed up the computation in our simulations by reducing calls to the integer optimization solver. However, note that the heuristic is merely that and cannot give a certificate for the non-existence of separation.

If the heuristic is unable to yield a separating cut, we implement the following integer optimization problem. Note

that $\langle \cdot, \cdot \rangle$ is the dot product of tensors obtained by flattening them into vectors.

$$\max_{\varphi, \theta} \langle c, \psi - \varphi \rangle$$
$$\begin{aligned}
\text{s.t. } \varphi_x &= \lambda y_{1,k} & x \in \mathcal{R} \\
y_{x,k} &\geq (-\theta_{x_k} - y_{x,k+1} - 1) & k \in [p-1], x \in \mathcal{R} \\
y_{x,k} &\geq (\theta_{x_k} + y_{x,k+1} - 1) & k \in [p-1], x \in \mathcal{R} \\
y_{x,k} &\leq (\theta_{x_k} - y_{x,k+1} + 1) & k \in [p-1], x \in \mathcal{R} \\
y_{x,k} &\leq (-\theta_{x_k} + y_{x,k+1} + 1) & k \in [p-1], x \in \mathcal{R} \\
y_{x,p} &= \theta_{x_p} & x \in \mathcal{R} \\
\theta_{x_k} &\in \{-1, 1\} & k \in [p], x \in \mathcal{R}
\end{aligned}$$
(9)

Since we are looking for weak separation, the integer optimization solver is made to terminate when a solution with an objective greater than Φ/K is found. In case of no such solution, the dual bound z from the solver serves as a no-separation certificate satisfying $\langle c, \psi - \varphi \rangle \leq z \leq \Phi$.

6. Numerical Experiments

In this section, we perform numerical experiments to demonstrate the efficacy and scalability of our tensor completion algorithm. The experiments were performed on computer server running a Linux operating system, with 16GB of RAM and an Intel Xeon Processor E5-2650L v3 (30M Cache, 1.80 GHz) that has 12 cores and became available in the year 2014. The algorithm was implemented in Python 3, and Gurobi v9.1 (Gurobi Optimization, LLC, 2021) was used for solving the integer programs (9).

We also performed experiments on benchmark algorithms in tensor completion, including the often-called 'workhorse' for numerical tensor problems, alternating least squares (ALS) (Kolda & Bader, 2009), and two state-of-the-art algorithms implemented in the PyTen package (Song et al., 2019), known as the simple low-rank tensor completion (SiLRTC) algorithm (Liu et al., 2012) and the trace norm regularized CP decomposition (TNCP) algorithm (Liu et al., 2014). PyTen is available at https://github.com/datamllab/pyten under a GPL 2 license.

The true tensor ψ is constructed in each experiment by taking a random convex combination of a random set of ten points from $\mathcal{S}_1$. This setup ensures that $\|\psi\|_\pm \leq 1$ and $\text{rank}(\psi) \leq 10$, which are provided as ground truth values during the experiments. Each experiment was performed with 100 repetitions. We recorded the normalized mean squared error (NMSE) to quantify the accuracy of tensor completion by each algorithm. NMSE is given by $\|\widehat{\psi} - \psi\|_F^2 / \|\psi\|_F^2$, which is a stricter measure than the error metric used in Corollary 5.2 because the statistical guarantee is not normalized.

Aiming to minimize the influence of hyper-parameter se-
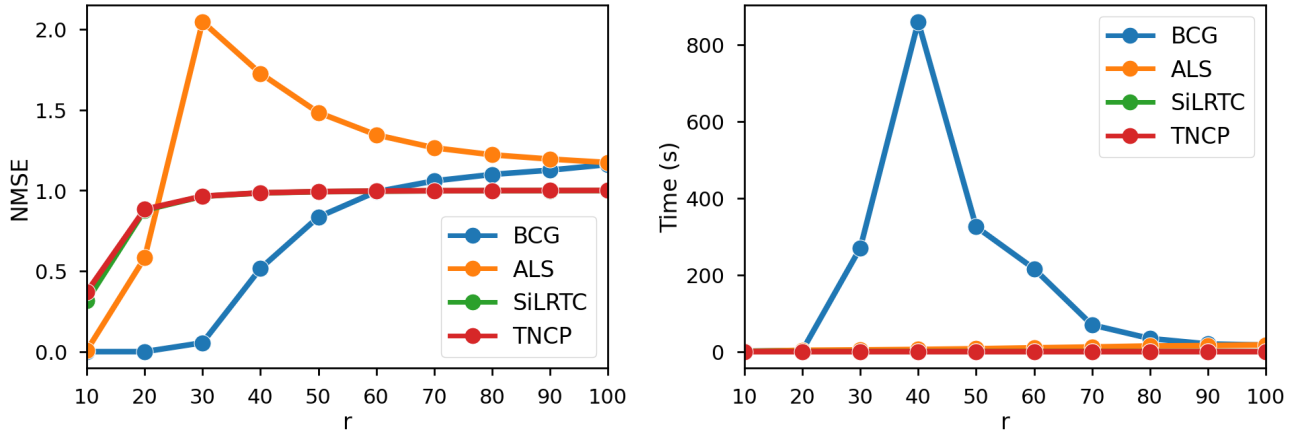
Figure 1. NMSE and computation time (in s) for order-3 tensors with size $r \times r \times r$ and $n = 1000$ samples.
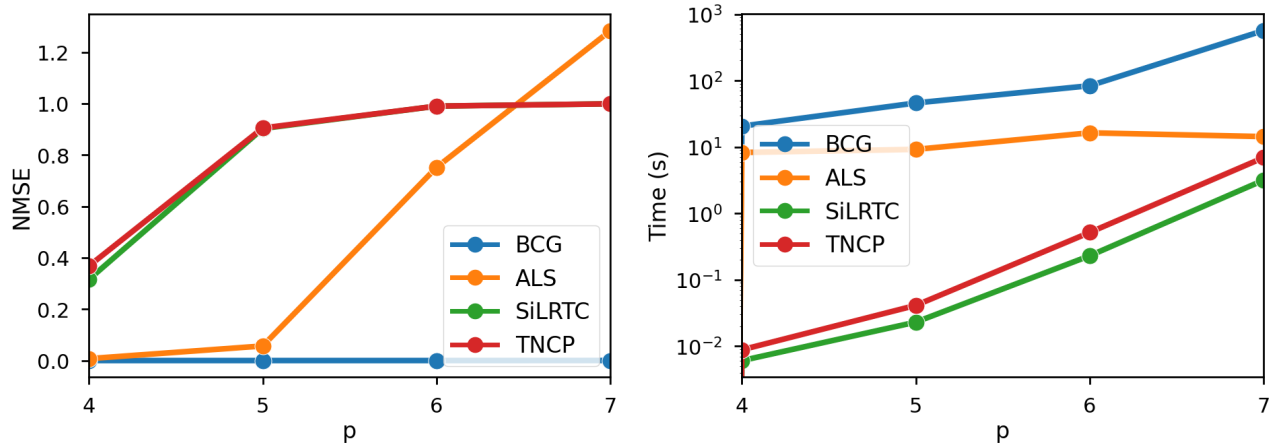

Figure 2. NMSE and computation time (in s) for increasing order tensors with size $10^{\times p}$ and $n = 10,000$ samples.
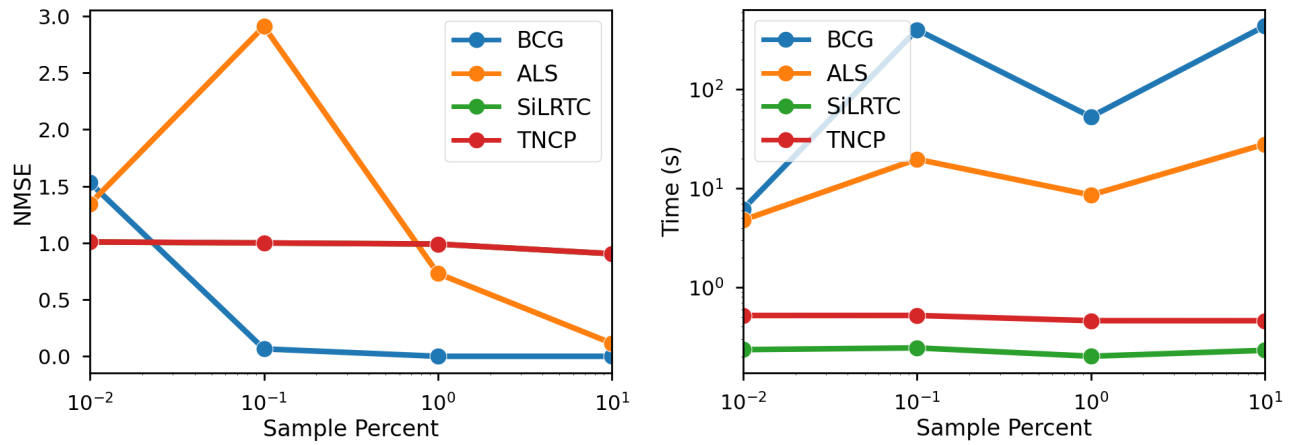

Figure 3. NMSE and computation time (in s) for tensors with size $10^{\times 6}$ and increasing n samples.
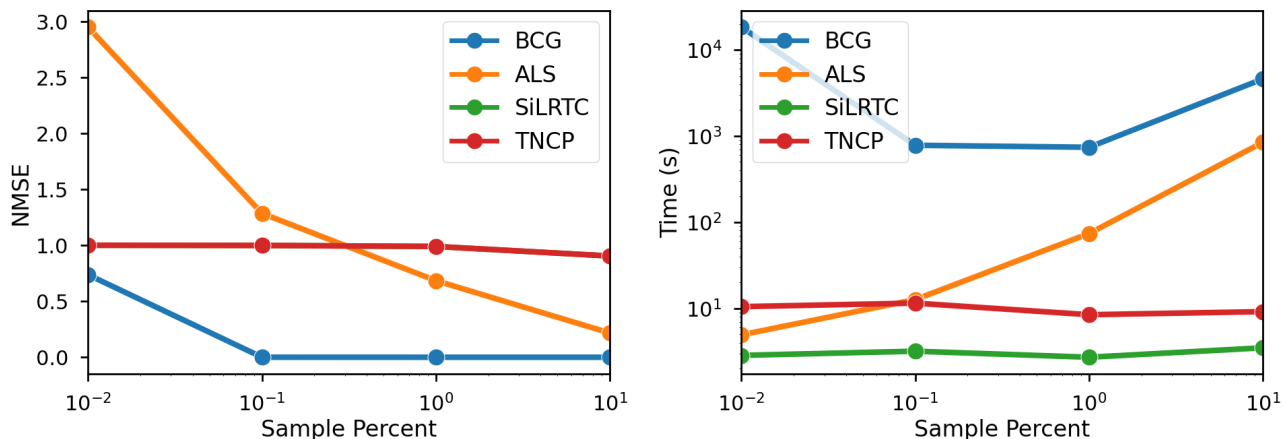
Figure 4. NMSE and computation time (in s) for tensors with size $10^{\times 7}$ and increasing n samples.

lection in the experiment results, we use the ground truth values during the numerical experiments when possible. In particular, constructing the true tensor ψ following the procedure as described above is useful for providing the ground truth values of λ in our algorithm and the value of k in ALS and TNCP. We also note that ALS tends to perform better under L2 regularization (Navasca et al., 2008), and thus we selected the L2 regularization hyperparameter of ALS such that it was favorable to the accuracy of ALS.

6.1. Increasing Tensor Dimension

The first set of results in Figure 1 is on tensors of order $p = 3$ with dimensions increasing from $r = 10$ to $r = 100$. In each experiment, we observe $n = 1000$ samples, allowing for repetition in indices. Figure 1 shows that our approach yields greater accuracy in a lower size of r while all algorithms do not perform at a satisfactory level (NMSE below 1) as r increases close to 100. Although our algorithm takes more computation time, it converges on the order of seconds to minutes for all dimensions. We note that TNCP and SiLRTC's NMSE values converge to 1 with increasing r as a naive solution, whose entries are equal to the average of all $y\langle i \rangle$, will lead to an NMSE of 1.

6.2. Increasing Tensor Order

The second set of results in Figure 2 is on tensors of increasing order p with dimension $r_i = 10$ for $i = 1, \ldots, p$. We observed $n = 10,000$ samples, again with indices sampled at random with replacement. This set of results shows that our algorithm achieves consistently higher accuracy as compared to the other methods while requiring more computation times. Yet, even for tensors with 10^7 entries, our algorithm is still able to converge within computation times on the order of minutes.

6.3. Increasing Sample Size

Our last set of results in Figures 3 and 4 is on tensors of sizes $10^{\times 6}$ and $10^{\times 7}$. In each experiment, the sample size is increased by one order of magnitude according to the values given in Figures 3 and 4 as the percentage of total entries, starting from 0.01% (using random sampling with replacement). The results demonstrate that our algorithm achieves considerably higher accuracy while requiring greater computation time. Nevertheless, our algorithm is able to converge within minutes for most cases, except for the case where the sample percent is $10^{-1}\%$ for a tensor of size 10^7 (i.e., approximately five hours).

7. Conclusion

We define a new tensor norm using the gauge of a specific polytope to develop an algorithm for (general) tensor completion. The algorithm successfully resolves the tension between practical computation and the information-theoretic rate: Our approach provably converges globally in a linear number of oracle calls while satisfying the information-theoretic sample complexity rate. Numerical experiments further demonstrated the efficacy and scalability of the algorithm. Next steps include efforts to further accelerate the algorithm, such that it can attain its high performance within computation times similar to benchmark algorithms.

Impact Statement

This paper presents work whose goal is to advance the field of Machine Learning. There are many potential societal consequences of our work, none of which we feel must be specifically highlighted here.

Acknowledgement

References

Acar, E., Dunlavy, D. M., Kolda, T. G., and Mørup, M. Scalable tensor factorizations for incomplete data. *Chemometrics and Intelligent Laboratory Systems*, 106(1):41–56, March 2011.

Aswani, A. Low-rank approximation and completion of positive tensors. *SIAM Journal on Matrix Analysis and Applications*, 37(3):1337–1364, 2016.

Barak, B. and Moitra, A. Noisy tensor completion via the sum-of-squares hierarchy. In *Conference on Learning Theory*, pp. 417–445. PMLR, 2016.

Bartlett, P. and Mendelson, S. Rademacher and gaussian complexities: Risk bounds and structural results. *J. Mach. Learn. Res.*, 2002.

Bazerque, J. A., Mateos, G., and Giannakis, G. B. Rank regularization and bayesian inference for tensor completion and extrapolation. *IEEE Transactions on Signal Processing*, 61(22):5689–5703, November 2013.

Braun, G., Pokutta, S., Tu, D., and Wright, S. Blended conditonal gradients. In *International Conference on Machine Learning*, pp. 735–743. PMLR, 2019.

Bugg, C., Chen, C., and Aswani, A. Nonnegative tensor completion via integer optimization. *Advances in Neural Information Processing Systems*, 35:10008–10020, 2022.

Cai, C., Li, G., Poor, H. V., and Chen, Y. Nonconvex low-rank tensor completion from noisy data. *Advances in neural information processing systems*, 32, 2019.

Chandrasekaran, V., Recht, B., Parrilo, P. A., and Willsky, A. S. The convex geometry of linear inverse problems. *Foundations of Computational mathematics*, 12(6):805–849, 2012.

Conforti, M., Cornuéjols, G., and Zambelli, G. *Integer Programming*, volume 271 of *Graduate Texts in Mathematics*. Springer International Publishing, Cham, 2014. ISBN 9783319110073 9783319110080.

de Silva, V. and Lim, L. Tensor rank and the ill-posedness of the best low-rank approximation problem. *SIAM J. Matrix Anal. Appl.*, 30(3):1084–1127, 2008. doi: 10.1137/06066518X. URL http://dx.doi.org/10.1137/06066518X.

Duarte, M. F. and Baraniuk, R. G. Kronecker compressive sensing. *IEEE Transactions on Image Processing*, 21(2):494–504, February 2012.

Friedland, S. and Lim, L.-H. The computational complexity of duality. *SIAM Journal on Optimization*, 26(4):2378–2393, 2016.

Gandy, S., Recht, B., and Yamada, I. Tensor completion and low-n-rank tensor recovery via convex optimization. *Inverse problems*, 27(2):025010, 2011.

Ge, H., Caverlee, J., and Lu, H. Taper: A contextual tensor-based approach for personalized expert recommendation. In *Proceedings of the 10th ACM Conference on Recommender Systems*, RecSys '16, pp. 261–268. Association for Computing Machinery, September 2016. ISBN 9781450340359.

Gurobi Optimization, LLC. Gurobi Optimizer Reference Manual, 2021. URL https://www.gurobi.com.

Hendrickx, J. M. and Olshevsky, A. Matrix p -norms are np-hard to approximate if $p \neq 1, 2, \infty$. *SIAM Journal on Matrix Analysis and Applications*, 31(5):2802–2812, 2010. doi: 10.1137/09076773X.

Hillar, C. and Lim, L.-H. Most tensor problems are np-hard. *J. ACM*, 60(6):45:1–45:39, November 2013. ISSN 0004-5411. doi: 10.1145/2512329. URL http://doi.acm.org/10.1145/2512329.

Jaggi, M. Revisiting frank-wolfe: Projection-free sparse convex optimization. In *Proceedings of the 30th International Conference on Machine Learning*, pp. 427–435, 2013.

Jain, S., Gutierrez, A., and Haupt, J. Noisy tensor completion for tensors with a sparse canonical polyadic factor. In *2017 IEEE International Symposium on Information Theory (ISIT)*, pp. 2153–2157, June 2017.

Javed, S., Bouwmans, T., and Jung, S. K. Stochastic decomposition into low rank and sparse tensor for robust background subtraction. In *6th International Conference on Imaging for Crime Prevention and Detection (ICDP-15)*, pp. 1–6, July 2015.

Karatzoglou, A., Amatriain, X., Baltrunas, L., and Oliver, N. Multiverse recommendation: n-dimensional tensor factorization for context-aware collaborative filtering. In *Proceedings of the fourth ACM conference on Recommender systems*, pp. 79–86. Association for Computing Machinery, September 2010. ISBN 9781605589060.

Kolda, T. and Bader, B. Tensor decompositions and applications. *SIAM Review*, 51(3):455–500, 2009.

Kressner, D., Steinlechner, M., and Vandereycken, B. Low-rank tensor completion by riemannian optimization. *BIT Numerical Mathematics*, 54(2):447–468, 2014.

Lecué, G. Empirical risk minimization is optimal for the convex aggregation problem. *Bernoulli*, 19(5B):2153–2166, 2013.

Li, X., Gunturk, B., and Zhang, L. Image demosaicing: A systematic survey. In *Visual Communications and Image Processing 2008*, volume 6822, pp. 68221J. International Society for Optics and Photonics, 2008.

Liu, J., Musialski, P., Wonka, P., and Ye, J. Tensor completion for estimating missing values in visual data. *IEEE transactions on pattern analysis and machine intelligence*, 35(1):208–220, 2012.

Liu, Y., Shang, F., Jiao, L., Cheng, J., and Cheng, H. Trace norm regularized candecomp/parafac decomposition with missing data. *IEEE transactions on cybernetics*, 45(11):2437–2448, 2014.

Montanari, A. and Sun, N. Spectral algorithms for tensor completion. *Communications on Pure and Applied Mathematics*, 71(11):2381–2425, 2018.

Navasca, C., De Lathauwer, L., and Kindermann, S. Swamp reducing technique for tensor decomposition. In *2008 16th European Signal Processing Conference*, pp. 1–5. IEEE, 2008.

Nemirovski, A. Topics in non-parametric statistics. *Ecole d'Eté de Probabilités de Saint-Flour*, 28:85, 2000.

Pan, W., Aswani, A., and Chen, C. Accelerated non-negative tensor completion via integer programming. *Frontiers in Applied Mathematics and Statistics*, 9, 2023. doi: 10.3389/fams.2023.1153184. URL https://www.frontiersin.org/articles/10.3389/fams.2023.1153184.

Qi, Y., Comon, P., and Lim, L.-H. Uniqueness of non-negative tensor approximations. *IEEE Transactions on Information Theory*, 62(4):2170–2183, 2016. doi: 10.1109/TIT.2016.2532906.

Rao, N., Shah, P., and Wright, S. Forward–backward greedy algorithms for atomic norm regularization. *IEEE Transactions on Signal Processing*, 63(21):5798–5811, 2015. doi: 10.1109/TSP.2015.2461515.

Rockafellar, R. and Wets, R. *Variational Analysis*. Springer, 2009.

Rockafellar, R. T. *Convex Analysis: (PMS-28)*. Princeton University Press, April 2015. ISBN 9781400873173.

Rohn, J. Computing the norm $\|a\|_{\infty,1}$ is np-hard. *Linear and Multilinear Algebra*, 47(3):195–204, 2000. doi: 10.1080/03081080008818644.

Signoretto, M., Van de Plas, R., De Moor, B., and Suykens, J. A. K. Tensor versus matrix completion: A comparison with application to spectral data. *IEEE Signal Processing Letters*, 18(7):403–406, July 2011.

Song, Q., Ge, H., Caverlee, J., and Hu, X. Tensor completion algorithms in big data analytics. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 13(1):1–48, 2019.

Srebro, N., Sridharan, K., and Tewari, A. Smoothness, low noise and fast rates. In *Advances in Neural Information Processing Systems*, pp. 2199–2207, 2010.

Tan, H., Cheng, B., Feng, J., Feng, G., Wang, W., and Zhang, Y.-J. Low-n-rank tensor recovery based on multi-linear augmented lagrange multiplier method. *Neurocomputing*, 119:144–152, November 2013.

Tomasi, G. and Bro, R. Parafac and missing values. *Chemometrics and Intelligent Laboratory Systems*, 75 (2):163–180, February 2005.

Tsybakov, A. B. Optimal rates of aggregation. In *Learning theory and kernel machines*, pp. 303–313. Springer, 2003.

Yuan, M. and Zhang, C.-H. On tensor completion via nuclear norm minimization. *Foundations of Computational Mathematics*, 16(4):1031–1068, 2016.

Zafarani, R., Abbasi, M. A., and Liu, H. *Social Media Mining: An Introduction*. Cambridge University Press, 1 edition, April 2014. ISBN 9781107018853 9781139088510.

Zhang, X., Wang, D., Zhou, Z., and Ma, Y. Robust low-rank tensor recovery with rectification and alignment. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(1):238–255, 2019.