

REWEIGHTED SOLUTIONS FOR WEIGHTED LOW RANK APPROXIMATION

Anonymous authors

Paper under double-blind review

ABSTRACT

Weighted low rank approximation is an important yet computationally challenging primitive with applications ranging from statistical analysis, model compression, and signal processing. To cope with the NP-hardness of this problem, prior work considers heuristics, bicriteria, or parameterized tractable algorithms to solve this problem. In this work, we introduce a new relaxed solution to weighted low rank approximation which outputs a matrix that is not necessarily low rank, but can be stored using very few parameters and gives provable approximation guarantees when the weight matrix has low rank. Our central idea is to use the weight matrix itself to reweight a low rank solution, which gives an extremely simple algorithm with remarkable empirical performance in applications to model compression. Our algorithm also gives nearly optimal communication complexity bounds for a natural distributed problem associated with this problem, for which we show matching communication lower bounds. Together, our communication complexity bounds show that the rank of the weight matrix provably parameterizes the communication complexity of weighted low rank approximation. We also obtain the first relative error guarantees for feature selection with a weighted objective.

1 INTRODUCTION

The approximation of matrices by one of lower rank has been, and continues to be, one of the most intensely studied and applied computational problems in statistics, machine learning, signal processing, and beyond. The classical approach to this problem is to approximate matrices $\mathbf{A} \in \mathbb{R}^{n \times d}$ by a rank k matrix $\tilde{\mathbf{A}} \in \mathbb{R}^{n \times d}$ that minimizes the Frobenius norm error

$$\|\mathbf{A} - \tilde{\mathbf{A}}\|_F^2 := \sum_{i=1}^n \sum_{j=1}^d |\mathbf{A}_{i,j} - \tilde{\mathbf{A}}_{i,j}|^2, \quad \text{rank}(\tilde{\mathbf{A}}) \leq k.$$

This problem is solved by the singular value decomposition (SVD), which can be computed in polynomial time. We will write \mathbf{A}_k to denote the optimal rank k approximation to \mathbf{A} in the Frobenius norm, and we will write $\mathbf{A}_{-k} := \mathbf{A} - \mathbf{A}_k$ to denote the error of this approximation.

While this simple choice often gives satisfactory results, this loss function treats all entries of the matrix uniformly when trying to fit $\tilde{\mathbf{A}}$, which may not exactly align with the practitioner’s desires if some of the entries are more crucial to fit than others. If one additionally has such information available in the form of nonnegative weights $\mathbf{W}_{i,j} \geq 0$ that reflects some measure of *importance* of each of the entries (i, j) , then this can be encoded in the loss function as follows:

$$\|\mathbf{A} - \tilde{\mathbf{A}}\|_{\mathbf{W},F}^2 := \sum_{i=1}^n \sum_{j=1}^d \mathbf{W}_{i,j}^2 \cdot |\mathbf{A}_{i,j} - \tilde{\mathbf{A}}_{i,j}|^2, \quad \text{rank}(\tilde{\mathbf{A}}) \leq k.$$

This problem is known as the *weighted low rank approximation (WLRA)* problem. We write $\mathbf{A} \circ \mathbf{B}$ to denote the entrywise product for two matrices \mathbf{A} and \mathbf{B} , so we may also write

$$\|\mathbf{A} - \tilde{\mathbf{A}}\|_{\mathbf{W},F}^2 := \|\mathbf{W} \circ (\mathbf{A} - \tilde{\mathbf{A}})\|_F^2 = \|\mathbf{W} \circ \mathbf{A} - \mathbf{W} \circ \tilde{\mathbf{A}}\|_F^2$$

The incorporation of weights into the low rank approximation problem gives this computational problem an incredible versatility for use in a long history of applications starting with its use in

factor analysis in the early statistical literature [Young \(1941\)](#). A popular special case is the *matrix completion* problem [Rennie & Srebro \(2005\)](#); [Candès & Tao \(2010\)](#); [Keshavan et al. \(2010\)](#), where the weights $\mathbf{W} \in \{0, 1\}^{n \times d}$ are binary and encode whether a given entry of \mathbf{A} is observed or not. This primitive has been useful in the design of recommender systems [Koren et al. \(2009\)](#); [Chen et al. \(2015\)](#); [Lee et al. \(2016\)](#), and has been famously applied in the 2006 Netflix Prize problem. More generally, the weights \mathbf{W} can be used to reflect the variance or number of samples obtained for each of the entries, so that more “uncertain” entries can influence the objective function less [Anandan & Irani \(2002\)](#); [Srebro & Jaakkola \(2003\)](#). In the past few years, weighted low rank approximation has also been used to improve model compression algorithms, especially those for large scale LLMs, based on low rank approximations of weight matrices by taking into account the importance of parameters [Arora et al. \(2016\)](#); [Hsu et al. \(2022\)](#); [Hua et al. \(2022\)](#). Given the rapid growth of large scale machine learning models, model compression techniques such as weighted low rank approximation are expected to bring high value to engineering efforts for these models. Other applications of weighted low rank approximation include ecology [Robin et al. \(2019\)](#); [Kidzinski et al. \(2022\)](#), background modeling [Li et al. \(2017\)](#); [Dutta et al. \(2018\)](#), computational biology [Tuzhilina et al. \(2022\)](#), and signal processing [Shpak \(1990\)](#); [Lu et al. \(1997\)](#).

Approximation algorithms have long been considered for efficient low rank approximation problems, and we formalize the approximation guarantee that we study in Definition [1.1](#).

Definition 1.1 (Approximate weighted low rank approximation). *Let $\mathbf{W} \in \mathbb{R}^{n \times d}$ be nonnegative, let $\mathbf{A} \in \mathbb{R}^{n \times d}$, and let $k \in \mathbb{N}$. Then in the κ -approximate rank k weighted low rank approximation problem, we seek to output a matrix $\tilde{\mathbf{A}} \in \mathbb{R}^{n \times d}$ such that*

$$\|\mathbf{A} - \tilde{\mathbf{A}}\|_{\mathbf{W}, F} \leq \kappa \min_{\text{rank}(\mathbf{A}') \leq k} \|\mathbf{A} - \mathbf{A}'\|_{\mathbf{W}, F}.$$

In Definition [1.1](#), we have purposefully under-specified requirements on $\tilde{\mathbf{A}}$. Part of this is to cope with the computational difficulty of WLRA. Indeed, while we ideally would like $\tilde{\mathbf{A}}$ to have rank at most k , solving for even an approximate such solution (with $\kappa = (1 + 1/\text{poly}(n))$) is an NP-hard problem [Gillis & Glineur \(2011\)](#). Furthermore, allowing for additional flexibility in the choice of $\tilde{\mathbf{A}}$ may still be useful as long as $\tilde{\mathbf{A}}$ satisfies some sort of “parameter reduction” guarantee. A common choice is to allow $\tilde{\mathbf{A}}$ to have rank $k' \geq k$ slightly larger than k , which is known as a *bicriteria* guarantee. In this work, we will show a new relaxation of the constraints on $\tilde{\mathbf{A}}$ that allows us to achieve new approximation guarantees for WLRA.

1.1 OUR RESULTS

We present our main contribution in Theorem [1.2](#), which gives a simple approach to WLRA, under the assumption that the weight matrix \mathbf{W} has low rank. We note that this assumption is very natural and captures natural cases, for example when \mathbf{W} has block structure, and has been motivated and studied in prior work [Razenshteyn et al. \(2016\)](#); [Ban et al. \(2019\)](#). We also empirically verify this assumption in our experiments. We defer a further discussion of the low rank \mathbf{W} assumption to Section [1.1.3](#) as well as prior works [Razenshteyn et al. \(2016\)](#); [Ban et al. \(2019\)](#).

The algorithm (shown in Algorithm [1](#)) that we propose is extremely simple: compute a rank rk approximation of $\mathbf{W} \circ \mathbf{A}$, and then divide the result entrywise by \mathbf{W} . Note that if we exactly compute the low rank approximation step by an SVD, then the optimal rank rk approximation $(\mathbf{W} \circ \mathbf{A})_{rk}$ given by the SVD requires only $(n + d)rk$ parameters to store, and \mathbf{W} also only requires nr parameters to store. Thus, the solution $\mathbf{W}^{\circ-1} \circ (\mathbf{W} \circ \mathbf{A})_{rk}$ can be stored in a total of $O((n + d)rk)$ parameters, which is nearly optimal for constant rank $r = O(1)$.

Algorithm 1 Weighted low rank approximation

input: input matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$, nonnegative weights $\mathbf{W} \in \mathbb{R}^{n \times d}$ with rank r .

output: approximate solution $\tilde{\mathbf{A}}$.

- 1: Compute a rank rk approximation $\tilde{\mathbf{A}}_{\mathbf{W}}$ of $\mathbf{W} \circ \mathbf{A}$
 - 2: **return** $\tilde{\mathbf{A}} := \mathbf{W}^{\circ-1} \circ \tilde{\mathbf{A}}_{\mathbf{W}}$
-

While our discussion thus far has simply used the SVD to compute the rank rk approximation $(\mathbf{W} \circ \mathbf{A})_{rk}$, we obtain other useful guarantees by allowing for approximate solutions $\tilde{\mathbf{A}}_{\mathbf{W}}$ that only approximately minimize $\|\mathbf{W} \circ \mathbf{A} - \tilde{\mathbf{A}}_{\mathbf{W}}\|_F$. For example, by computing the rank rk approximation $\tilde{\mathbf{A}}_{\mathbf{W}}$ using faster randomized approximation algorithms for the SVD [Clarkson & Woodruff (2013); Musco & Musco (2015); Avron et al. (2017)], we obtain algorithms for WLRA with similar running time. In general, we prove the following theorem:

Theorem 1.2. *Let $\mathbf{W} \in \mathbb{R}^{n \times d}$ be a nonnegative weight matrix with rank r . Let $\mathbf{A} \in \mathbb{R}^{n \times d}$ and let $k \in \mathbb{N}$. Suppose that $\tilde{\mathbf{A}}_{\mathbf{W}} \in \mathbb{R}^{n \times d}$ satisfies*

$$\|\mathbf{W} \circ \mathbf{A} - \tilde{\mathbf{A}}_{\mathbf{W}}\|_F^2 \leq \kappa \min_{\text{rank}(\mathbf{A}') \leq rk} \|\mathbf{W} \circ \mathbf{A} - \mathbf{A}'\|_F^2 = \kappa \|(\mathbf{W} \circ \mathbf{A})_{-rk}\|_F^2$$

and let $\tilde{\mathbf{A}} := \mathbf{W}^{\circ-1} \circ \tilde{\mathbf{A}}_{\mathbf{W}}$, where $\mathbf{W}^{\circ-1} \in \mathbb{R}^{n \times d}$ denotes the entrywise inverse of \mathbf{W} . Then,

$$\|\mathbf{A} - \tilde{\mathbf{A}}\|_{\mathbf{W},F}^2 \leq \kappa \min_{\text{rank}(\mathbf{A}') \leq k} \|\mathbf{A} - \mathbf{A}'\|_{\mathbf{W},F}^2$$

In particular, we obtain a solution with $\kappa = (1 + \varepsilon)$ in running time $O(\text{nnz}(\mathbf{A})) + \tilde{O}(n(rk)^2/\varepsilon + \text{poly}(rk/\varepsilon))$ by using randomized low rank approximation algorithms of [Avron et al. (2017)].

We prove Theorem 1.2 in Section 2.

We note that as stated, the approximation given by Algorithm 1 may not always be desirable, since in general, $\mathbf{W}^{\circ-1}$ cannot be computed without multiplying out the low rank factors of \mathbf{W} . However, we show in Lemma A.1 that for a broad family of structured matrices formed by the sum of support-disjoint rank 1 matrices and a sparse matrix, $\mathbf{W}^{\circ-1}$ can in fact be stored and applied in the same time as \mathbf{W} . These capture a large number of commonly used weight matrix patterns in practice, such as Low-Rank Plus Sparse, Low-Rank Plus Diagonal, Low-Rank Plus Block Diagonal, Monotone Missing Data Pattern, and Low-Rank Plus Banded matrices [Musco et al. (2021)].

For general matrices, we present a more practical alternative in Algorithm 2 where we compute a low rank approximation of $\mathbf{W}^{\circ-1}$, so that the entrywise inverse can be applied efficiently.

Algorithm 2 Weighted low rank approximation (practical)

input: input matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$, nonnegative weights $\mathbf{W} \in \mathbb{R}^{n \times d}$ with rank r .

output: approximate solution $\tilde{\mathbf{A}}$.

- 1: Compute a rank rk approximation $\tilde{\mathbf{W}}_{\text{inv}}$ of $\mathbf{W}^{\circ-1}$
 - 2: Compute a rank rk approximation $\tilde{\mathbf{A}}_{\mathbf{W}}$ of $\tilde{\mathbf{W}}_{\text{inv}}^{\circ-1} \circ \mathbf{A}$
 - 3: **return** $\tilde{\mathbf{A}} := \tilde{\mathbf{W}}_{\text{inv}} \circ \tilde{\mathbf{A}}_{\mathbf{W}}$
-

1.1.1 COLUMN SUBSET SELECTION FOR WEIGHTED LOW RANK APPROXIMATION

Another advantage of allowing for approximation algorithms for computing low rank approximations to $\mathbf{W} \circ \mathbf{A}$ is that we can employ *column subset selection* approaches to low rank approximation [Frieze et al. (2004); Deshpande & Vempala (2006); Drineas et al. (2006, 2008); Boutsidis et al. (2016); Altschuler et al. (2016)]. That is, it is known that the Frobenius norm low rank approximation problem admits $(1 + \varepsilon)$ -approximate low rank approximations whose left factor is formed by a subset of at most $O(k/\varepsilon)$ columns of the input matrix. In particular, these results show the existence of approximate solutions to the low rank approximation problem that preserves the sparsity of the input matrix, and thus can lead to a reduced solution size when the input matrix has sparse columns. Furthermore, column subset selection solutions to low rank approximation give a natural approach for *unsupervised feature selection*. Thus, as a corollary of Theorem 1.2, we obtain the first relative error guarantee for unsupervised feature selection with a weighted Frobenius norm objective. Weaker additive error guarantees were previously studied by [Dai (2023); Axiotis & Yasuda (2023)]¹.

Corollary 1.3 (Column subset selection for weighted low rank approximation). *There is an algorithm that computes a subset $S \subseteq [d]$ of $|S| = O(rk/\varepsilon)$ columns and $\mathbf{X} \in \mathbb{R}^{|S| \times d}$ such that*

$$\|\mathbf{A} - \mathbf{W}^{\circ-1} \circ ((\mathbf{W} \circ \mathbf{A})^S \mathbf{X})\|_{\mathbf{W},F}^2 \leq (1 + \varepsilon) \min_{\text{rank}(\mathbf{A}') \leq k} \|\mathbf{A} - \mathbf{A}'\|_{\mathbf{W},F}^2$$

¹The result of [Dai (2023)] contained an error, which we correct, tighten, and simplify in Appendix D.

where for a matrix $\mathbf{B} \in \mathbb{R}^{n \times d}$, $\mathbf{B}|^S$ denotes the matrix formed by the columns of \mathbf{B} indexed by S .

Proof. This follows from Theorem 1.2 by computing the rank rk approximation $\tilde{\mathbf{A}}_{\mathbf{W}}$ to $\mathbf{W} \circ \mathbf{A}$ via column subset selection algorithms given by, e.g., Boutsidis et al. (2016). \square

Note that in Corollary 1.3 the approximation $\mathbf{W}^{\circ-1} \circ ((\mathbf{W} \circ \mathbf{A})|^S \mathbf{X})$ only depends on \mathbf{A} through the columns $\mathbf{A}|^S$, and thus giving an approach to column subset selection with a weighted objective.

1.1.2 NEARLY OPTIMAL COMMUNICATION COMPLEXITY BOUNDS

As a consequence of Corollary 1.3, we obtain another important result for WLRA in the setting of *communication complexity*. Here, we obtain nearly optimal communication complexity bounds for constant factor approximations (i.e. $\kappa = O(1)$) to distributed WLRA for a wide range of parameters. While many works have studied distributed LRA in depth (Sarlós, 2006; Clarkson & Woodruff, 2009; 2013; Macua et al., 2010; Kannan et al., 2014; Ghashami et al., 2016; Boutsidis et al., 2016), we are surprisingly the first to initiate a study of this problem for WLRA.

The communication setting we consider is as follows. We have two players, Alice and Bob, where Alice has an input matrix \mathbf{A} and would like to communicate an approximate WLRA solution to Bob. Communication complexity is of great interest in modern computing, where exchanging bits can be a critical bottleneck in large scale computation. While we consider two players in this discussion for simplicity, our algorithms also apply to a distributed computing setting, where the columns of the input matrix are partitioned among m servers as m matrices $\mathbf{A}^{(1)}, \mathbf{A}^{(2)}, \dots, \mathbf{A}^{(m)}$, and some central coordinator outputs a WLRA of the concatenation $\mathbf{A} = [\mathbf{A}^{(1)}, \mathbf{A}^{(2)}, \dots, \mathbf{A}^{(m)}]$ of these columns.

Definition 1.4 (WLRA: communication game). *Let Alice and Bob be two players. Let $\mathbf{W} \in \mathbb{Z}^{n \times d}$ be nonnegative, let $\mathbf{A} \in \mathbb{Z}^{n \times d}$, and let $k \in \mathbb{N}$. Furthermore, let \mathbf{W} and \mathbf{A} have entries at most $\mathbf{W}_{i,j}, |\mathbf{A}_{i,j}| \leq \text{poly}(nd)$. We let both Alice and Bob receive the weight matrix \mathbf{W} as input, and we give only Alice the input matrix \mathbf{A} . We say that Alice and Bob solve the κ -approximate rank k weighted low rank approximation communication game using B bits of communication if Alice sends at most B bits to Bob, and Bob outputs any matrix $\tilde{\mathbf{A}} \in \mathbb{R}^{n \times d}$ satisfying*

$$\|\mathbf{A} - \tilde{\mathbf{A}}\|_{\mathbf{W},F} \leq \kappa \min_{\text{rank}(\mathbf{A}') \leq k} \|\mathbf{A} - \mathbf{A}'\|_{\mathbf{W},F}.$$

Suppose that \mathbf{A} has columns which each have at most s nonzero entries. Then, the solution given by Corollary 1.3 can be communicated to Bob using just $O(srk/\varepsilon + rkd)$ numbers ($O(srk/\varepsilon)$ for the $O(rk/\varepsilon)$ columns of \mathbf{A} and $O(rkd)$ for \mathbf{X}), or $O((srk/\varepsilon + rkd) \log(nd))$ bits under our bit complexity assumptions. Thus, when the number of columns d is at most the column sparsity s , then we obtain an algorithm which uses only $O((srk/\varepsilon) \log(nd))$ bits of communication. More generally, if the columns of \mathbf{A} are distributed among m servers, then a solution can be computed using $O((msrk/\varepsilon) \log(nd))$ bits of communication by using work of Boutsidis et al. (2016).

In fact, we show a nearly matching communication lower bound. In particular, we show that $\Omega(srk)$ bits of communication is required to output *any* matrix (not necessarily structured) that achieves a weighted Frobenius norm loss that is *any* finite factor within the optimal solution. Our lower bound is information theoretic, and also immediately implies an $\Omega(msrk)$ bit lower bound in the distributed setting of m servers if each server must output a solution, as considered by Boutsidis et al. (2016).

Theorem 1.5. *Let \mathbf{W} be a binary block diagonal mask (Definition 3.2) and let $k \in \mathbb{N}$. Suppose that a randomized algorithm solves, for every $\mathbf{C} \in \mathbb{Z}^{n \times n}$ with at most s nonzero entries in each column, the κ -approximate weighted low rank approximation problem using B bits of communication with probability at least $2/3$, for any $1 \leq \kappa < \infty$. If $s, k \leq n/r$, then $B = \Omega(srk)$.*

By proving a nearly tight communication complexity bound of $\tilde{\Theta}(srk)$ for computing constant factor WLRA, we arrive at the following qualitative observation: *the rank r of the weight matrix \mathbf{W} parameterizes the communication complexity of WLRA*. Note that a similar conclusion was drawn for the *computational* complexity of WLRA in the work of Razenshteyn et al. (2016), where it was shown that WLRA is fixed parameter tractable in the parameter r , and also must have running time exponential in r under natural complexity theoretic assumptions. Thus, we believe that one important contribution of our work is to provide further evidence, both empirical and theoretical, that the rank r of the weight matrix \mathbf{W} is a natural parameter to consider when studying WLRA.

1.1.3 EXPERIMENTS

We demonstrate the empirical performance of our WLRA algorithms through experiments for model compression tasks. This application of WLRA was suggested by [Hsu et al. \(2022\)](#); [Hua et al. \(2022\)](#), which we find to be a particularly relevant application of weighted low rank approximation due to the trend of extremely large models. In the model compression setting, we wish to approximate the hidden layer weight matrices of neural networks by much smaller matrices. A classical way to do this is to use low rank approximation ([Sainath et al. 2013](#); [Kim et al. 2016](#); [Chen et al. 2018](#)). While this often gives reasonable results, the works of [Hsu et al. \(2022\)](#); [Hua et al. \(2022\)](#) show that significant improvements can be obtained by taking into account the importance of each of the parameters in the LRA problem. We thus conduct our experiments in this setting.

We first show in Section [4.1](#) that the importance matrices arising this application are indeed very low rank. We may interpret this phenomenon intuitively: we hypothesize that the importance score of some parameter $\mathbf{A}_{i,j}$ is essentially the product of the importance of the corresponding input i and the importance of the corresponding output j . This observation may be of independent interest, and also empirically justifies the low rank weight matrix assumption that we make in this work, as well as works of [Razenshteyn et al. \(2016\)](#); [Ban et al. \(2019\)](#). While WLRA with a rank 1 weight matrix is known to be solvable efficiently via the SVD, our result shows that general low rank weight matrices also yield efficient algorithms via the SVD.

Next in Section [4.2](#), we conduct experiments which demonstrates the superiority of our methods in practice. Of the algorithms that we compare to, an expectation-minimization (EM) approach of [Srebro & Jaakkola \(2003\)](#) gives the smallest loss albeit with a very high running time, and our algorithm nearly matches this loss with an order of magnitude lower running time. We also show that this solution can be refined with EM, producing the best trade-off between efficiency and accuracy. One of the baselines we compare is a sampling algorithm of [Dai \(2023\)](#), whose analysis contains an error which we correct, simplify, and tighten.

1.2 RELATED WORK

We survey a number of related works on approximation algorithms for weighted low rank approximation. One of the earliest algorithms for this problem is a natural EM approach proposed by [Srebro & Jaakkola \(2003\)](#). Another related approach is to parameterize the low rank approximation \mathbf{A} as the product $\mathbf{U}\mathbf{V}$ of two matrices $\mathbf{U} \in \mathbb{R}^{n \times k}$ and $\mathbf{V} \in \mathbb{R}^{k \times d}$ and alternately minimize the two matrices, known as *alternating least squares*. This algorithm has been studied in a number of works ([Hastie et al. 2015](#); [Li et al. 2016](#); [Song et al. 2023](#)). The work of [Bhaskara et al. \(2021\)](#) proposes an approach to weighted low rank approximation based on a greedy pursuit, where rank one factors are iteratively added based on an SVD of the gradient matrix. Finally, fixed parameter tractable algorithms have been considered in [Razenshteyn et al. \(2016\)](#); [Ban et al. \(2019\)](#) based on sketching techniques.

2 APPROXIMATION ALGORITHMS

The following simple observation is the key idea behind Theorem [1.2](#):

Lemma 2.1. *Let $\mathbf{W}, \mathbf{A}' \in \mathbb{R}^{n \times d}$ with $\text{rank}(\mathbf{W}) \leq r$ and $\text{rank}(\mathbf{A}') \leq k$. Then, $\text{rank}(\mathbf{W} \circ \mathbf{A}') \leq rk$.*

Proof. Since $\text{rank}(\mathbf{W}) \leq k$, it can be written as $\mathbf{W} = \sum_{i=1}^r \mathbf{u}_i \mathbf{v}_i^\top$ for $\mathbf{u}_i \in \mathbb{R}^n$ and $\mathbf{v}_i \in \mathbb{R}^d$. Then,

$$\mathbf{W} \circ \mathbf{A}' = \sum_{i=1}^r (\mathbf{u}_i \mathbf{v}_i^\top) \circ \mathbf{A}' = \sum_{i=1}^r \text{diag}(\mathbf{u}_i) \mathbf{A}' \text{diag}(\mathbf{v}_i)$$

so $\mathbf{W} \circ \mathbf{A}'$ is the sum of r matrices, each of which is rank k . Thus, $\mathbf{W} \circ \mathbf{A}'$ has rank at most rk . \square

Using Lemma [2.1](#), we obtain the following:

Theorem 1.2. *Let $\mathbf{W} \in \mathbb{R}^{n \times d}$ be a nonnegative weight matrix with rank r . Let $\mathbf{A} \in \mathbb{R}^{n \times d}$ and let $k \in \mathbb{N}$. Suppose that $\tilde{\mathbf{A}}_{\mathbf{W}} \in \mathbb{R}^{n \times d}$ satisfies*

$$\|\mathbf{W} \circ \mathbf{A} - \tilde{\mathbf{A}}_{\mathbf{W}}\|_F^2 \leq \kappa \min_{\text{rank}(\mathbf{A}') \leq rk} \|\mathbf{W} \circ \mathbf{A} - \mathbf{A}'\|_F^2 = \kappa \|(\mathbf{W} \circ \mathbf{A})_{-rk}\|_F^2$$

and let $\tilde{\mathbf{A}} := \mathbf{W}^{\circ-1} \circ \tilde{\mathbf{A}}_{\mathbf{W}}$, where $\mathbf{W}^{\circ-1} \in \mathbb{R}^{n \times d}$ denotes the entrywise inverse of \mathbf{W} . Then,

$$\|\mathbf{A} - \tilde{\mathbf{A}}\|_{\mathbf{W},F}^2 \leq \kappa \min_{\text{rank}(\mathbf{A}') \leq k} \|\mathbf{A} - \mathbf{A}'\|_{\mathbf{W},F}^2$$

In particular, we obtain a solution with $\kappa = (1 + \varepsilon)$ in running time $O(\text{nnz}(\mathbf{A})) + \tilde{O}(n(rk)^2/\varepsilon + \text{poly}(rk/\varepsilon))$ by using randomized low rank approximation algorithms of [Avron et al. \(2017\)](#).

Proof. Note that $\|\mathbf{W}^{\circ-1} \circ \tilde{\mathbf{A}}_{\mathbf{W}} - \mathbf{A}\|_{\mathbf{W},F}^2 = \|\tilde{\mathbf{A}}_{\mathbf{W}} - \mathbf{W} \circ \mathbf{A}\|_F^2$, which is at most $\kappa \|(\mathbf{W} \circ \mathbf{A})_{-rk}\|_F^2$ by assumption. On the other hand for any rank k matrix \mathbf{A}' , $\|\mathbf{A}' - \mathbf{A}\|_{\mathbf{W},F} = \|\mathbf{W} \circ \mathbf{A}' - \mathbf{W} \circ \mathbf{A}\|_F$ can be lower bounded by $\|(\mathbf{W} \circ \mathbf{A})_{-rk}\|_F$ since $\mathbf{W} \circ \mathbf{A}'$ has rank at most rk by Lemma [2.1](#). Thus,

$$\|\mathbf{W}^{\circ-1} \circ \tilde{\mathbf{A}}_{\mathbf{W}} - \mathbf{A}\|_{\mathbf{W},F}^2 \leq \kappa \|(\mathbf{W} \circ \mathbf{A})_{-rk}\|_F^2 \leq \kappa \min_{\text{rank}(\mathbf{A}') \leq k} \|\mathbf{A} - \mathbf{A}'\|_{\mathbf{W},F}^2. \quad \square$$

3 COMMUNICATION COMPLEXITY BOUNDS

We show that our approach to weighted low rank approximation in Theorem [1.2](#) gives nearly optimal bounds for this problem in the setting of communication complexity.

Our first result is an upper bound for the communication game in Definition [1.4](#)

Theorem 3.1. *Let $\mathbf{W} \in \mathbb{Z}^{n \times d}$ be a nonnegative rank k weight matrix and let $\mathbf{A} \in \mathbb{Z}^{n \times d}$ be an input matrix with at most s nonzero entries in each column. There is an algorithm which solves the $(1 + \varepsilon)$ -approximate weighted low rank approximation communication game (Definition [1.4](#)) using at most $B = O((srk/\varepsilon + rkd) \log(nd))$ bits of communication.*

Proof. The algorithm is simply to use the column subset selection-based WLRA algorithm of Corollary [1.3](#) and then to send the columns of \mathbf{A} indexed by the column subset S and \mathbf{X} . \square

On the other hand, we show a communication complexity lower bound showing that the number of bits B exchanged by Alice and Bob must be at least $\Omega(rsk)$. Our lower bound holds even when the weight matrix \mathbf{W} is the following simple binary matrix.

Definition 3.2 (Block diagonal mask). *Let $r \in \mathbb{N}$ and let n be an integer multiple of r . Then, $\mathbf{W} \in \{0, 1\}^{n \times n}$ is the block diagonal mask associated with these parameters if \mathbf{W} is the $r \times r$ block diagonal matrix with diagonal blocks given by the $n/r \times n/r$ all ones matrix and off-diagonal blocks given by the $n/r \times n/r$ all zeros matrix.*

We give our communication complexity lower bound in the following theorem.

Theorem 1.5. *Let \mathbf{W} be a binary block diagonal mask (Definition [3.2](#)) and let $k \in \mathbb{N}$. Suppose that a randomized algorithm solves, for every $\mathbf{C} \in \mathbb{Z}^{n \times n}$ with at most s nonzero entries in each column, the κ -approximate weighted low rank approximation problem using B bits of communication with probability at least $2/3$, for any $1 \leq \kappa < \infty$. If $s, k \leq n/r$, then $B = \Omega(srk)$.*

Proof. Let $\mathbf{A}_{\text{dense}} \in \{0, 1\}^{sr \times k}$ be a uniformly random binary matrix, and let $\mathbf{A}_{\text{pad}} \in \{0, 1\}^{n \times n/r}$ be formed by padding the columns of $\mathbf{A}_{\text{dense}}$ with $n/r - k$ columns of all zeros and padding each block of s contiguous rows with $n/r - s$ rows of all zeros. For $j \in [r]$, let $\mathbf{A}_{\text{pad}}^{(j)}$ denote the restriction of \mathbf{A}_{pad} to the j th contiguous block of n/r rows. We then construct $\mathbf{A} \in \mathbb{R}^{n \times n}$ by horizontally concatenating r copies of \mathbf{A}_{pad} .

Note that an optimal rank k approximation can achieve 0 loss in the \mathbf{W} -weighted Frobenius norm. Indeed, we can take \mathbf{A}^* to be the horizontal concatenation of r copies of \mathbf{A}_{pad} . Since \mathbf{A}_{pad} has rank k , \mathbf{A}^* also has rank k . Furthermore, on the j -th nonzero blocks of \mathbf{W} , \mathbf{A}_{pad} has the same entries as $\mathbf{A}_{\text{pad}}^{(j)}$. Thus, it follows that an approximation $\tilde{\mathbf{A}}$ that achieves any finite approximation factor κ must exactly recover \mathbf{A} , restricted to the support of \mathbf{W} . In turn, this means that such an approximation $\tilde{\mathbf{A}}$ can also be used to recover $\mathbf{A}_{\text{dense}}$.

It now follows by a standard information theoretic argument that $B = \Omega(srk)$ (see Appendix [B](#) for further details). \square

4 EXPERIMENTS

As discussed in Section 1.1.3, we conduct experiments for weighted low rank approximation in the setting of model compression (as proposed by Hsu et al. (2022); Hua et al. (2022)). In our experiments, we train a basic multilayer perceptron (MLP) on four image datasets, `mnist`, `fashion_mnist`, `smallnorb`, and `colorectal_histology` which were selected from the `tensorflow_datasets` catalogue for simplicity of processing (e.g. fixed feature size, no need for embeddings, etc). We then compute a matrix of importances of each of the parameters in a hidden layer of the MLP given by the Fisher information matrix. Finally, we compute a weighted low rank approximation of the hidden layer matrix using the Fisher information matrix as the weights \mathbf{W} .

Our experiments are conducted on a 2019 MacBook Pro with a 2.6 GHz 6-Core Intel Core i7 processor. All code used in the experiments are available in the supplement.

Table 1: Datasets used in experiments

Dataset	Image dim.	Flattened dim.	Neurons	Matrix dim.
<code>mnist</code>	(28, 28, 1)	784	128	784×128
<code>fashion_mnist</code>	(28, 28, 1)	784	128	784×128
<code>smallnorb</code>	(96, 96, 1)	9216	1024	9216×1024
<code>colorectal_histology</code>	(150, 150, 3)	67500	1024	67500×1024

4.1 THE LOW RANK WEIGHT MATRIX ASSUMPTION IN PRACTICE

We first demonstrate that for the task of model compression, the weight matrix is approximately low rank in practice. The weight matrix \mathbf{W} in this setting is the empirical Fisher information matrix of the hidden layer weights \mathbf{A} , where the empirical Fisher information of the (i, j) -th entry $\mathbf{A}_{i,j}$ is given by

$$\mathbf{W}_{i,j} := \mathbb{E}_{\mathbf{x} \sim \mathcal{D}} \left[\left(\frac{\partial}{\partial \mathbf{A}_{i,j}} \mathcal{L}(\mathbf{x}; \mathbf{A}) \right)^2 \right]$$

where $\mathcal{L}(\mathbf{x}; \mathbf{A})$ denotes the loss of the neural network on the data point \mathbf{x} and hidden layer weights \mathbf{A} , and \mathcal{D} denotes the empirical distribution (that is, the uniform distribution over the training data).

Plots of the empirical Fisher matrix (Figure 1) reveal nontrivial low rank structure to the matrices, and the spectrum of the Fisher matrix confirms that the vast majority of the Frobenius norm is contained in the first singular value (Table 2). We also plot the spectrum itself in Figure 5 in the appendix.



Figure 1: Low rank structure of Fisher weight matrices

Table 2: Singular value distribution of Fisher matrix

Dataset	% Frobenius mass in first singular value
<code>mnist</code>	95.4%
<code>fashion_mnist</code>	95.9%
<code>smallnorb</code>	99.9%
<code>colorectal_histology</code>	99.3%

4.2 APPROXIMATION QUALITY AND RUNNING TIME

In this section, we compare the performance of our Algorithm 2 (denoted as `svd_w` in the following discussion) with a variety of previously proposed algorithms for weighted low rank approximation.

We consider the following algorithms: `adam`, `em`, `greedy`, `sample`, and `svd`, which we next explain in detail. We first consider `adam`, in which we simply parameterize the WLRA problem as an optimization problem in the factorized representation \mathbf{UV} for factors $\mathbf{U} \in \mathbb{R}^{n \times k}$ and $\mathbf{V} \in \mathbb{R}^{k \times d}$, and optimize this loss function using the Adam optimizer provided in the `tensorflow` library. Such an approach is well-studied for the standard low rank approximation problem Li et al. (2018); Ye & Du (2021), and empirically performs well for weighted low rank approximation as well. This was run for 100 epochs, with an initial learning rate of 1.0 decayed by a factor of 0.7 every 10 steps. The `em` algorithm was proposed by Srebro & Jaakkola (2003) for the WLRA problem, and involves iteratively “filling in” missing values and recomputing a low rank approximation. In the experiments, we run 25 iterations. The `greedy` algorithm is a greedy basis pursuit algorithm proposed by Bhaskara et al. (2021) and iteratively adds new directions to the low rank approximation by taking an SVD of the gradient of the objective. Similar algorithms were also studied in Shalev-Shwartz et al. (2011); Khanna et al. (2017); Axiotis & Sviridenko (2021) for general rank-constrained convex optimization problems. The `sample` algorithm is a row norm sampling approach studied by Dai (2023). Finally, `svd` simply computes an SVD of the original matrix \mathbf{W} , without regard to the weights \mathbf{W} .

We compute low rank approximations for ranks 1 through 20 on four datasets, and plot the loss and the running time against the rank in Figures 2 and 3, respectively. The values in the figures are tabulated at ranks 20, 10, and 5 in Tables 3, 4, and 5 in the supplement. We observe that our `svd_w` algorithm performs among the best in the approximation loss (Figure 2), nearly matching the approximation quality achieved by much more computationally expensive algorithms such as `adam` and `em`, while requiring much less computational time (Figure 3).

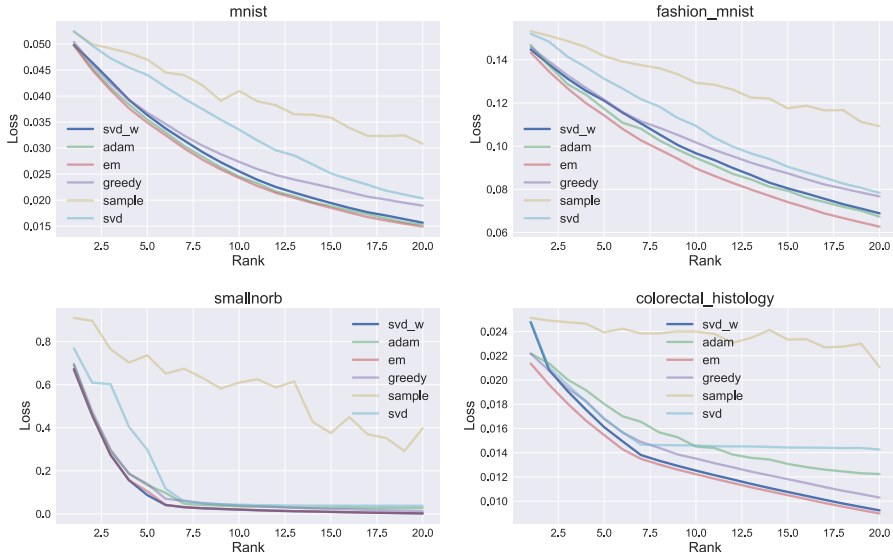


Figure 2: Fisher-weighted low rank approximation loss of weighted low rank approximation algorithms for model compression four datasets. Results are averaged over 5 trials.

While in some cases the `em` algorithm may eventually produce a better solution, we note that our `svd_w` may be improved by initializing the `em` algorithm with this solution, which produces an algorithm which quickly produces a superior solution with much fewer iterations (Figure 4).

5 CONCLUSION

In this work, we studied new algorithms for the weighted low rank approximation problem, which has countless applications in statistics, machine learning, and signal processing. We propose an approach

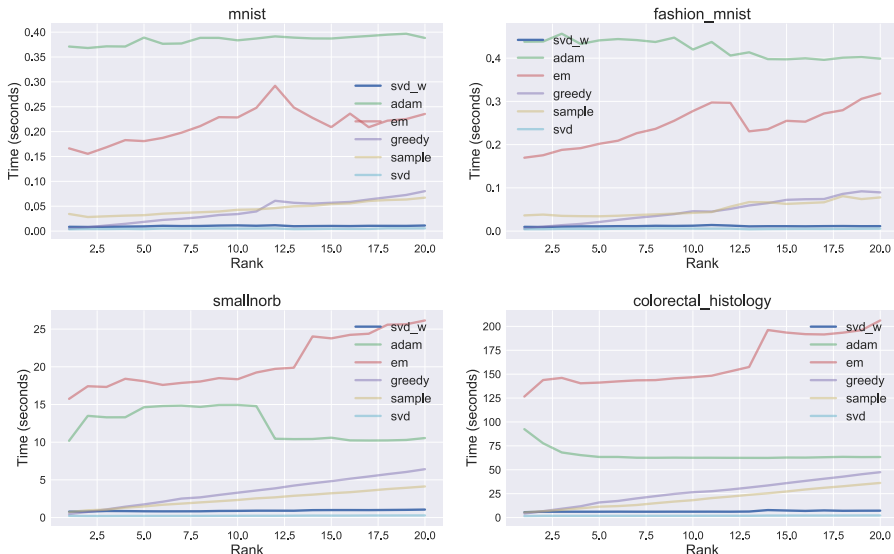


Figure 3: Running time of weighted low rank approximation algorithms for model compression four datasets. Results are averaged over 5 trials.

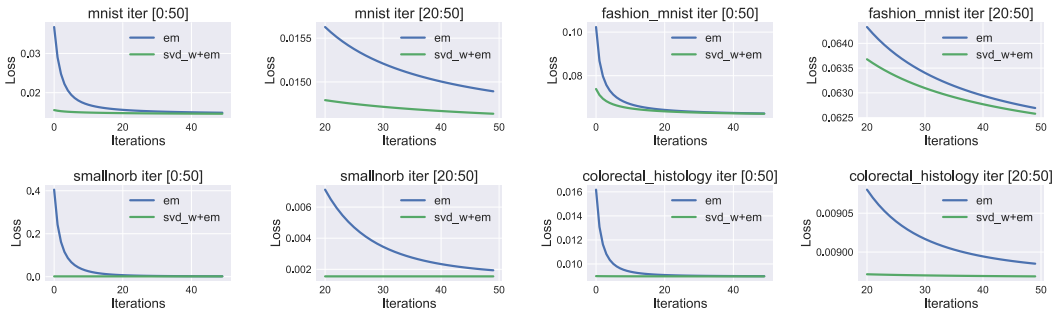


Figure 4: Improving the `svd_w` solution with `em` iterations for a rank 20 approximation.

based on reweighting a low rank matrix, which is a novel class of relaxed solutions to the WLRA problem, and give provable guarantees under the assumption that the weight matrix \mathbf{W} has low rank. Theoretically, this allows us to obtain an algorithm for WLRA with nearly optimal communication complexity, for which we show nearly matching communication complexity lower bounds, which shows that the rank of the weight matrix tightly parameterizes the communication complexity of this problem. We also give the first guarantees for column subset selection for weighted low rank approximation, which gives a notion of feature selection with a weighted objective. Finally, we show that in practice, our approach gives a highly efficient algorithm that outperforms prior algorithms for WLRA, particularly when combined with refinement using expectation-maximization.

BIBLIOGRAPHY

Jason M. Altschuler, Aditya Bhaskara, Gang Fu, Vahab S. Mirrokni, Afshin Rostamizadeh, and Morteza Zadimoghaddam. Greedy column subset selection: New bounds and distributed algorithms. In Maria-Florina Balcan and Kilian Q. Weinberger (eds.), *Proceedings of the 33rd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19-24, 2016*, volume 48 of *JMLR Workshop and Conference Proceedings*, pp. 2539–2548. JMLR.org, 2016. URL <http://proceedings.mlr.press/v48/altschuler16.html>

P. Anandan and Michal Irani. Factorization with uncertainty. *Int. J. Comput. Vis.*, 49(2-3): 101–116, 2002. doi: 10.1023/A:1020137420717. URL <https://doi.org/10.1023/A:>

1020137420717

- Sanjeev Arora, Yuanzhi Li, Yingyu Liang, Tengyu Ma, and Andrej Risteski. A latent variable model approach to pmi-based word embeddings. *Trans. Assoc. Comput. Linguistics*, 4:385–399, 2016. doi: 10.1162/tacl_a_00106. URL https://doi.org/10.1162/tacl_a_00106.
- Haim Avron, Kenneth L. Clarkson, and David P. Woodruff. Sharper bounds for regularized data fitting. In Klaus Jansen, José D. P. Rolim, David Williamson, and Santosh S. Vempala (eds.), *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2017, August 16-18, 2017, Berkeley, CA, USA*, volume 81 of *LIPICs*, pp. 27:1–27:22. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017. doi: 10.4230/LIPICs.APPROX-RANDOM.2017.27. URL <https://doi.org/10.4230/LIPICs.APPROX-RANDOM.2017.27>.
- Kyriakos Axiotis and Maxim Sviridenko. Local search algorithms for rank-constrained convex optimization. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021. URL https://openreview.net/forum?id=tH6_VWZjoo.
- Kyriakos Axiotis and Taisuke Yasuda. Performance of ℓ_1 regularization for sparse convex optimization. *CoRR*, abs/2307.07405, 2023. doi: 10.48550/arXiv.2307.07405. URL <https://doi.org/10.48550/arXiv.2307.07405>.
- Frank Ban, David P. Woodruff, and Qiuyi (Richard) Zhang. Regularized weighted low rank approximation. In Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d’Alché-Buc, Emily B. Fox, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pp. 4061–4071, 2019. URL <https://proceedings.neurips.cc/paper/2019/hash/90aef91f0d9e7c3be322bd7bae41617d-Abstract.html>.
- Aditya Bhaskara, Aravinda Kanchana Ruwanpathirana, and Maheshakya Wijewardena. Additive error guarantees for weighted low rank approximation. In Marina Meila and Tong Zhang (eds.), *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pp. 874–883. PMLR, 2021. URL <http://proceedings.mlr.press/v139/bhaskara21a.html>.
- Christos Boutsidis, David P. Woodruff, and Peilin Zhong. Optimal principal component analysis in distributed and streaming models. In Daniel Wichs and Yishay Mansour (eds.), *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016, Cambridge, MA, USA, June 18-21, 2016*, pp. 236–249. ACM, 2016. doi: 10.1145/2897518.2897646. URL <https://doi.org/10.1145/2897518.2897646>.
- Emmanuel J. Candès and Terence Tao. The power of convex relaxation: near-optimal matrix completion. *IEEE Trans. Inf. Theory*, 56(5):2053–2080, 2010. doi: 10.1109/TIT.2010.2044061. URL <https://doi.org/10.1109/TIT.2010.2044061>.
- Chao Chen, Dongsheng Li, Yingying Zhao, Qin Lv, and Li Shang. WEMAREC: accurate and scalable recommendation through weighted and ensemble matrix approximation. In Ricardo Baeza-Yates, Mounia Lalmas, Alistair Moffat, and Berthier A. Ribeiro-Neto (eds.), *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval, Santiago, Chile, August 9-13, 2015*, pp. 303–312. ACM, 2015. doi: 10.1145/2766462.2767718. URL <https://doi.org/10.1145/2766462.2767718>.
- Patrick H. Chen, Si Si, Yang Li, Ciprian Chelba, and Cho-Jui Hsieh. Groupreduce: Block-wise low-rank approximation for neural language model shrinking. In Samy Bengio, Hanna M. Wallach, Hugo Larochelle, Kristen Grauman, Nicolò Cesa-Bianchi, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pp. 11011–11021, 2018. URL <https://proceedings.neurips.cc/paper/2018/hash/a2b8a85a29b2d64ad6f47275bf1360c6-Abstract.html>.

- Kenneth L. Clarkson and David P. Woodruff. Numerical linear algebra in the streaming model. In Michael Mitzenmacher (ed.), *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009, Bethesda, MD, USA, May 31 - June 2, 2009*, pp. 205–214. ACM, 2009. doi: 10.1145/1536414.1536445. URL <https://doi.org/10.1145/1536414.1536445>.
- Kenneth L. Clarkson and David P. Woodruff. Low rank approximation and regression in input sparsity time. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum (eds.), *Symposium on Theory of Computing Conference, STOC'13, Palo Alto, CA, USA, June 1-4, 2013*, pp. 81–90. ACM, 2013. doi: 10.1145/2488608.2488620. URL <https://doi.org/10.1145/2488608.2488620>.
- Thomas M. Cover and Joy A. Thomas. *Elements of information theory (2. ed.)*. Wiley, 2006. ISBN 978-0-471-24195-9. URL <http://www.elementsofinformationtheory.com/>.
- Yucheng Dai. On algorithms for weighted low rank approximation. Master’s thesis, Carnegie Mellon University, 2023.
- Amit Deshpande and Santosh S. Vempala. Adaptive sampling and fast low-rank matrix approximation. In Josep Díaz, Klaus Jansen, José D. P. Rolim, and Uri Zwick (eds.), *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, 9th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems, APPROX 2006 and 10th International Workshop on Randomization and Computation, RANDOM 2006, Barcelona, Spain, August 28-30 2006, Proceedings*, volume 4110 of *Lecture Notes in Computer Science*, pp. 292–303. Springer, 2006. doi: 10.1007/11830924_28. URL https://doi.org/10.1007/11830924_28.
- Petros Drineas, Michael W. Mahoney, and S. Muthukrishnan. Subspace sampling and relative-error matrix approximation: Column-row-based methods. In Yossi Azar and Thomas Erlebach (eds.), *Algorithms - ESA 2006, 14th Annual European Symposium, Zurich, Switzerland, September 11-13, 2006, Proceedings*, volume 4168 of *Lecture Notes in Computer Science*, pp. 304–314. Springer, 2006. doi: 10.1007/11841036_29. URL https://doi.org/10.1007/11841036_29.
- Petros Drineas, Michael W. Mahoney, and S. Muthukrishnan. Relative-error CUR matrix decompositions. *SIAM J. Matrix Anal. Appl.*, 30(2):844–881, 2008. doi: 10.1137/07070471X. URL <https://doi.org/10.1137/07070471X>.
- Aritra Dutta, Xin Li, and Peter Richtárik. Weighted low-rank approximation of matrices and background modeling. *CoRR*, abs/1804.06252, 2018. URL <http://arxiv.org/abs/1804.06252>.
- Alan M. Frieze, Ravi Kannan, and Santosh S. Vempala. Fast monte-carlo algorithms for finding low-rank approximations. *J. ACM*, 51(6):1025–1041, 2004. doi: 10.1145/1039488.1039494. URL <https://doi.org/10.1145/1039488.1039494>.
- Mina Ghashami, Edo Liberty, Jeff M. Phillips, and David P. Woodruff. Frequent directions: Simple and deterministic matrix sketching. *SIAM J. Comput.*, 45(5):1762–1792, 2016. doi: 10.1137/15M1009718. URL <https://doi.org/10.1137/15M1009718>.
- Nicolas Gillis and François Glineur. Low-rank matrix approximation with weights or missing data is NP-hard. *SIAM J. Matrix Anal. Appl.*, 32(4):1149–1165, 2011. doi: 10.1137/110820361. URL <https://doi.org/10.1137/110820361>.
- Trevor Hastie, Rahul Mazumder, Jason D. Lee, and Reza Zadeh. Matrix completion and low-rank SVD via fast alternating least squares. *J. Mach. Learn. Res.*, 16:3367–3402, 2015. doi: 10.5555/2789272.2912106. URL <https://dl.acm.org/doi/10.5555/2789272.2912106>.
- Yen-Chang Hsu, Ting Hua, Sungen Chang, Qian Lou, Yilin Shen, and Hongxia Jin. Language model compression with weighted low-rank factorization. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net, 2022. URL <https://openreview.net/forum?id=uPv9Y3qmAI5>.
- Ting Hua, Yen-Chang Hsu, Felicity Wang, Qian Lou, Yilin Shen, and Hongxia Jin. Numerical optimizations for weighted low-rank estimation on language models. In Yoav Goldberg, Zornitsa Kozareva, and Yue Zhang (eds.), *Proceedings of the 2022 Conference on Empirical Methods in*

- Natural Language Processing, EMNLP 2022, Abu Dhabi, United Arab Emirates, December 7-11, 2022*, pp. 1404–1416. Association for Computational Linguistics, 2022. doi: 10.18653/v1/2022.emnlp-main.91. URL <https://doi.org/10.18653/v1/2022.emnlp-main.91>.
- Ravi Kannan, Santosh S. Vempala, and David P. Woodruff. Principal component analysis and higher correlations for distributed data. In Maria-Florina Balcan, Vitaly Feldman, and Csaba Szepesvári (eds.), *Proceedings of The 27th Conference on Learning Theory, COLT 2014, Barcelona, Spain, June 13-15, 2014*, volume 35 of *JMLR Workshop and Conference Proceedings*, pp. 1040–1057. JMLR.org, 2014. URL <http://proceedings.mlr.press/v35/kannan14.html>.
- Raghunandan H. Keshavan, Andrea Montanari, and Sewoong Oh. Matrix completion from a few entries. *IEEE Trans. Inf. Theory*, 56(6):2980–2998, 2010. doi: 10.1109/TIT.2010.2046205. URL <https://doi.org/10.1109/TIT.2010.2046205>.
- Rajiv Khanna, Ethan R. Elenberg, Alexandros G. Dimakis, Joydeep Ghosh, and Sahand N. Negahban. On approximation guarantees for greedy low rank optimization. In Doina Precup and Yee Whye Teh (eds.), *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, volume 70 of *Proceedings of Machine Learning Research*, pp. 1837–1846. PMLR, 2017. URL <http://proceedings.mlr.press/v70/khanna17a.html>.
- Lukasz Kidzinski, Francis K. C. Hui, David I. Warton, and Trevor J. Hastie. Generalized matrix factorization: efficient algorithms for fitting generalized linear latent variable models to large data arrays. *J. Mach. Learn. Res.*, 23:291:1–291:29, 2022. URL <http://jmlr.org/papers/v23/20-1104.html>.
- Yong-Deok Kim, Eunhyeok Park, Sungjoo Yoo, Taelim Choi, Lu Yang, and Dongjun Shin. Compression of deep convolutional neural networks for fast and low power mobile applications. In Yoshua Bengio and Yann LeCun (eds.), *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*, 2016. URL <http://arxiv.org/abs/1511.06530>.
- Yehuda Koren, Robert M. Bell, and Chris Volinsky. Matrix factorization techniques for recommender systems. *Computer*, 42(8):30–37, 2009. doi: 10.1109/MC.2009.263. URL <https://doi.org/10.1109/MC.2009.263>.
- Joonseok Lee, Seungyeon Kim, Guy Lebanon, Yoram Singer, and Samy Bengio. LLORMA: local low-rank matrix approximation. *J. Mach. Learn. Res.*, 17:15:1–15:24, 2016. URL <http://jmlr.org/papers/v17/14-301.html>.
- Xin Li, Aritra Dutta, and Peter Richtárik. A batch-incremental video background estimation model using weighted low-rank approximation of matrices. In *2017 IEEE International Conference on Computer Vision Workshops, ICCV Workshops 2017, Venice, Italy, October 22-29, 2017*, pp. 1835–1843. IEEE Computer Society, 2017. doi: 10.1109/ICCVW.2017.217. URL <https://doi.org/10.1109/ICCVW.2017.217>.
- Yuanzhi Li, Yingyu Liang, and Andrej Risteski. Recovery guarantee of weighted low-rank approximation via alternating minimization. In Maria-Florina Balcan and Kilian Q. Weinberger (eds.), *Proceedings of the 33rd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19-24, 2016*, volume 48 of *JMLR Workshop and Conference Proceedings*, pp. 2358–2367. JMLR.org, 2016. URL <http://proceedings.mlr.press/v48/li16.html>.
- Yuanzhi Li, Tengyu Ma, and Hongyang Zhang. Algorithmic regularization in over-parameterized matrix sensing and neural networks with quadratic activations. In Sébastien Bubeck, Vianney Perchet, and Philippe Rigollet (eds.), *Conference on Learning Theory, COLT 2018, Stockholm, Sweden, 6-9 July 2018*, volume 75 of *Proceedings of Machine Learning Research*, pp. 2–47. PMLR, 2018. URL <http://proceedings.mlr.press/v75/li18a.html>.
- W-S Lu, S-C Pei, and P-H Wang. Weighted low-rank approximation of general complex matrices and its application in the design of 2-d digital filters. *IEEE Transactions on Circuits and Systems I: Fundamental Theory and Applications*, 44(7):650–655, 1997.

- Sergio Valcarcel Macua, Pavle Belanovic, and Santiago Zazo. Consensus-based distributed principal component analysis in wireless sensor networks. In *2010 IEEE 11th International Workshop on Signal Processing Advances in Wireless Communications (SPAWC)*, pp. 1–5. IEEE, 2010.
- Cameron Musco and Christopher Musco. Randomized block krylov methods for stronger and faster approximate singular value decomposition. In Corinna Cortes, Neil D. Lawrence, Daniel D. Lee, Masashi Sugiyama, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pp. 1396–1404, 2015. URL <https://proceedings.neurips.cc/paper/2015/hash/1efa39bcaec6f3900149160693694536-Abstract.html>
- Cameron Musco, Christopher Musco, and David P. Woodruff. Simple heuristics yield provable algorithms for masked low-rank approximation. In James R. Lee (ed.), *12th Innovations in Theoretical Computer Science Conference, ITCS 2021, January 6-8, 2021, Virtual Conference*, volume 185 of *LIPIcs*, pp. 6:1–6:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021. doi: 10.4230/LIPIcs.ITCS.2021.6. URL <https://doi.org/10.4230/LIPIcs.ITCS.2021.6>.
- Ilya P. Razenshteyn, Zhao Song, and David P. Woodruff. Weighted low rank approximations with provable guarantees. In Daniel Wichs and Yishay Mansour (eds.), *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016, Cambridge, MA, USA, June 18-21, 2016*, pp. 250–263. ACM, 2016. doi: 10.1145/2897518.2897639. URL <https://doi.org/10.1145/2897518.2897639>.
- Jason D. M. Rennie and Nathan Srebro. Fast maximum margin matrix factorization for collaborative prediction. In Luc De Raedt and Stefan Wrobel (eds.), *Machine Learning, Proceedings of the Twenty-Second International Conference (ICML 2005), Bonn, Germany, August 7-11, 2005*, volume 119 of *ACM International Conference Proceeding Series*, pp. 713–719. ACM, 2005. doi: 10.1145/1102351.1102441. URL <https://doi.org/10.1145/1102351.1102441>
- Geneviève Robin, Julie Josse, Eric Moulines, and Sylvain Sardy. Low-rank model with covariates for count data with missing values. *J. Multivar. Anal.*, 173:416–434, 2019. doi: 10.1016/j.jmva.2019.04.004. URL <https://doi.org/10.1016/j.jmva.2019.04.004>.
- Tara N. Sainath, Brian Kingsbury, Vikas Sindhwani, Ebru Arisoy, and Bhuvana Ramabhadran. Low-rank matrix factorization for deep neural network training with high-dimensional output targets. In *IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2013, Vancouver, BC, Canada, May 26-31, 2013*, pp. 6655–6659. IEEE, 2013. doi: 10.1109/ICASSP.2013.6638949. URL <https://doi.org/10.1109/ICASSP.2013.6638949>.
- Tamás Sarlós. Improved approximation algorithms for large matrices via random projections. In *47th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2006), 21-24 October 2006, Berkeley, California, USA, Proceedings*, pp. 143–152. IEEE Computer Society, 2006. doi: 10.1109/FOCS.2006.37. URL <https://doi.org/10.1109/FOCS.2006.37>
- Shai Shalev-Shwartz, Alon Gonen, and Ohad Shamir. Large-scale convex minimization with a low-rank constraint. In Lise Getoor and Tobias Scheffer (eds.), *Proceedings of the 28th International Conference on Machine Learning, ICML 2011, Bellevue, Washington, USA, June 28 - July 2, 2011*, pp. 329–336. Omnipress, 2011. URL https://icml.cc/2011/papers/232_icmlpaper.pdf
- Dale J Shpak. A weighted-least-squares matrix decomposition method with application to the design of two-dimensional digital filters. In *Proceedings of the 33rd Midwest Symposium on Circuits and Systems*, pp. 1070–1073. IEEE, 1990.
- Zhao Song, Mingquan Ye, Junze Yin, and Lichen Zhang. Efficient alternating minimization with applications to weighted low rank approximation. *CoRR*, abs/2306.04169, 2023. doi: 10.48550/arXiv.2306.04169. URL <https://doi.org/10.48550/arXiv.2306.04169>.
- Nathan Srebro and Tommi S. Jaakkola. Weighted low-rank approximations. In Tom Fawcett and Nina Mishra (eds.), *Machine Learning, Proceedings of the Twentieth International Conference*

(*ICML 2003*), August 21-24, 2003, Washington, DC, USA, pp. 720–727. AAAI Press, 2003. URL <http://www.aaai.org/Library/ICML/2003/icml03-094.php>.

Elena Tuzhilina, Trevor J. Hastie, and Mark R. Segal. Principal curve approaches for inferring 3D chromatin architecture. *Biostatistics*, 23(2):626–642, 2022. ISSN 1465-4644. doi: 10.1093/biostatistics/kxaa046. URL <https://doi.org/10.1093/biostatistics/kxaa046>.

Tian Ye and Simon S. Du. Global convergence of gradient descent for asymmetric low-rank matrix factorization. In Marc’Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan (eds.), *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pp. 1429–1439, 2021. URL <https://proceedings.neurips.cc/paper/2021/hash/0af854284f4ab0cfea8fcfd889cbb41a-Abstract.html>.

Gale Young. Maximum likelihood estimation and factor analysis. *Psychometrika*, 6:49–53, 1941. ISSN 0033-3123. doi: 10.1007/BF02288574. URL <https://doi.org/10.1007/BF02288574>.