

Efficient Optimization Methods for Extreme Similarity Learning with Nonlinear Embeddings

Bowen Yuan
National Taiwan University
f03944049@csie.ntu.edu.tw

Pengrui Quan
University of California, Los Angeles
prquan@g.ucla.edu

Yu-Sheng Li
National Taiwan University
r07922087@csie.ntu.edu.tw

Chih-Jen Lin
National Taiwan University
cjlin@csie.ntu.edu.tw

ABSTRACT

We study the problem of learning similarity by using nonlinear embedding models (e.g., neural networks) from all possible pairs. This problem is well-known for its difficulty of training with the extreme number of pairs. For the special case of using linear embeddings, many studies have addressed this issue of handling all pairs by considering certain loss functions and developing efficient optimization algorithms. This paper aims to extend results for general nonlinear embeddings. First, we finish detailed derivations and provide clean formulations for efficiently calculating some building blocks of optimization algorithms such as function, gradient evaluation, and Hessian-vector product. The result enables the use of many optimization methods for extreme similarity learning with nonlinear embeddings. Second, we study some optimization methods in detail. Due to the use of nonlinear embeddings, implementation issues different from linear cases are addressed. In the end, some methods are shown to be highly efficient for extreme similarity learning with nonlinear embeddings.

CCS CONCEPTS

• Computing methodologies → Supervised learning.

KEYWORDS

Similarity learning, Representation learning, Non-convex optimization, Newton methods, Neural networks

ACM Reference Format:

Bowen Yuan, Yu-Sheng Li, Pengrui Quan, and Chih-Jen Lin. 2021. Efficient Optimization Methods for Extreme Similarity Learning with Nonlinear Embeddings. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '21), August 14–18, 2021, Virtual Event, Singapore*. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3447548.3467363>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
KDD '21, August 14–18, 2021, Virtual Event, Singapore

© 2021 Association for Computing Machinery.
ACM ISBN 978-1-4503-8332-5/21/08...\$15.00
<https://doi.org/10.1145/3447548.3467363>

1 INTRODUCTION

Many applications can be cast in the problem of learning similarity between a pair of two entities referred to as the left and the right entities respectively. For example, in recommender systems, the similarity of a user-item pair indicates the preference of the user on the item. In search engines, the similarity of a query-document pair can be used as the relevance between the query and the document. In multi-label classifications, for any instance-label pair with high similarity, the instance can be categorized to the label.

A popular approach for similarity learning is to train an embedding model as the representation of each entity in the pair, such that any pair with high similarity are mapped to two close vectors in the embedding space, and vice versa. For the choice of the embedding model, some recent works [30] report the superiority of nonlinear over conventional linear ones. A typical example of applying nonlinear embedding models is the two-tower structure illustrated in Figure 1, where a multi-layer neural network serves as the embedding model of each entity. Some successful uses in real-world applications have been reported [7, 11, 12, 23, 24].

While many works only consider observed pairs for similarity learning, more and more works argue that better performance can be achieved by considering all possible pairs. For example, recommender systems with implicit feedback face a one-class scenario, where all observed pairs are labeled as similar while the dissimilar ones are missing. To achieve better similarity learning, a widely-used setting [16, 26, 27] is to include all unobserved pairs as dissimilar ones. Another example is counterfactual learning [22, 28], where because the observed pairs carry the selection bias caused by confounders, an effective way to eliminate the bias is to additionally consider the unobserved pairs by imputing their labels.

However, for many real-world scenarios, both the number m of left entities and the number n of right entities can be tremendous. The learning procedure is challenging as a prohibitive $O(mn)$ cost occurs by directly applying any optimization algorithm. To avoid the $O(mn)$ complexity, one can subsample unobserved pairs, but it has been shown that such subsampled settings are inferior to the non-sampled setting [17, 25]. We refer to the problem of the similarity learning from extremely large mn pairs as *extreme similarity learning*.

To tackle the high $O(mn)$ complexity, the aforementioned works that can handle all pairs consider a loss function with a certain structure for the unobserved pairs. Then they are able to replace the $O(mn)$ complexity with a much smaller $O(m+n)$ one. In particular, the function value, gradient, or other information can be calculated




Figure 1: An illustration of the two-tower structure, adopted from [18].

in $O(m+n)$ cost, so various types of optimization methods can be considered. Most exiting works consider linear embeddings (e.g., matrix factorization in [9, 16, 25]), where the optimization problem is often in a multi-block convex form. Thus many consider a block-wise setting to sequentially minimize convex sub-problems. The needed gradient or other information on each block of variables can also be calculated in $O(m+n)$.

If general nonlinear embeddings are considered, so far few works have studied the optimization algorithm. This work aims to fill the gap with the following main contributions.

- To calculate function value, gradient, or other information in $O(m+n)$ cost, the extension from cases in linear embeddings seems to be possible, but detailed derivations have not been available. We finish tedious calculations and provide clean formulations. This result enables the use of many optimization methods for extreme similarity learning with nonlinear embeddings.
- We then study some optimization methods in detail. Due to the use of nonlinear embeddings, some implementation issues different from linear cases are addressed. In the end, some methods are shown to be highly efficient for extreme similarity learning with nonlinear embeddings.

The paper is organized as follows. A review on extreme similarity learning is in Section 2. We derive efficient computation for some important components in Section 3. In Section 4, we demonstrate that many optimization algorithms can be used for extreme similarity learning with nonlinear embeddings. Section 5 describes some implementation issues. Finally, we present experiments on large-scale data sets in Section 6, and conclude in Section 7. Table 1 gives main notations in this paper. The supplementary materials and data/code for experiments are available at https://www.csie.ntu.edu.tw/~cjlin/papers/similarity_learning/.

2 EXTREME SIMILARITY LEARNING

We review extreme similarity learning problems.

2.1 Problem Formulation

Many real-world applications can be modeled as a learning problem with an incomplete similarity matrix \mathbf{R} of m left entities and n right entities. The set of observed pairs is denoted as $\mathbb{O} \subseteq \{1, \dots, m\} \times \{1, \dots, n\}$ where typically $|\mathbb{O}| \ll mn$. Only $R_{ij}, \forall (i, j) \in \mathbb{O}$ are

Table 1: Main notation

Notation	Description
(i, j)	(left entity i , right entity j) pair
m, n	numbers of left entities and right entities
\mathbf{u}, \mathbf{v}	feature vectors of a left entity and a right entity
\mathbb{O}	set of observed pairs
$y(\cdot), \theta \in \mathbb{R}^{D_u+D_v}$	a similarity function and its variables
$L(\theta)$	objective function
\mathbf{R}	observed incomplete similarity matrix
$\hat{\mathbf{Y}}$	similarity matrix predicted by $y(\cdot)$
$\tilde{\mathbf{Y}}$	imputed similarity matrix (for unobserved pairs)
$f(\cdot), g(\cdot)$	embedding models of left and right entities
$\mathbf{p} \in \mathbb{R}^k, \mathbf{q} \in \mathbb{R}^k$	embedding vectors of left and right entities
$F(f), F(g)$	cost of operations related to $f(\cdot)$, and $g(\cdot)$

revealed. Besides \mathbf{R} , we assume that side features $\mathbb{U} = \{\dots, \mathbf{u}_i, \dots\}$ of left entities and $\mathbb{V} = \{\dots, \mathbf{v}_j, \dots\}$ of right entities are available. Our goal is to find a similarity function $y: \mathbb{R}^D \rightarrow \mathbb{R}$,

$$\hat{Y}_{ij} = y(\theta; \mathbf{u}_i, \mathbf{v}_j), \quad (1)$$

where $\theta = \begin{bmatrix} \theta^u \\ \theta^v \end{bmatrix}$ is the vectorized variables of y with $\theta^u \in \mathbb{R}^{D_u}$, $\theta^v \in \mathbb{R}^{D_v}$, and $D = D_u + D_v$. For the similarity function, we consider the following dot product similarity that has been justified in a recent comparative study of similarity functions in recommender systems [18]

$$y(\theta; \mathbf{u}_i, \mathbf{v}_j) = f(\theta^u; \mathbf{u}_i)^\top g(\theta^v; \mathbf{v}_j),$$

where $f: \mathbb{R}^{D_u} \rightarrow \mathbb{R}^k$ and $g: \mathbb{R}^{D_v} \rightarrow \mathbb{R}^k$ are two embedding models that learn representations of left and right entities, respectively; see Figure 1. In this work, we focus on nonlinear embedding models. For convenience, we abbreviate $f(\theta^u; \mathbf{u}_i)$ and $g(\theta^v; \mathbf{v}_j)$ to $f^i(\theta^u)$ and $g^j(\theta^v)$, respectively. We also write the embedding vectors as

$$\mathbf{p}_i = f^i(\theta^u), \mathbf{q}_j = g^j(\theta^v) \quad (2)$$

so that

$$\hat{Y}_{ij} = \mathbf{p}_i^\top \mathbf{q}_j = f^i(\theta^u)^\top g^j(\theta^v). \quad (3)$$

By considering all mn pairs, we learn θ through solving

$$\min_{\theta} L(\theta), \quad (4)$$

where the objective function is

$$L(\theta) = \sum_{i=1}^m \sum_{j=1}^n \ell_{ij}(Y_{ij}, \hat{Y}_{ij}) + \lambda \Psi(\theta), \quad (5)$$

ℓ_{ij} is an entry-wise twice-differentiable loss function convex in \hat{Y}_{ij} , $\Psi(\theta)$ is a twice-differentiable and strongly convex regularizer to avoid overfitting, and $\lambda > 0$ is the corresponding parameter. We assume Ψ is simple so that in all complexity analyses we ignore the costs related to $\Psi(\theta)$, $\nabla \Psi(\theta)$, and $\nabla^2 \Psi(\theta)$. In (5), we further assume

$$Y_{ij} = \begin{cases} R_{ij} & (i, j) \in \mathbb{O}, \\ \tilde{Y}_{ij} & (i, j) \notin \mathbb{O}, \end{cases} \quad (6)$$

where for any unobserved pair $(i, j) \notin \mathbb{O}$ we impute \tilde{Y}_{ij} as an artificial label of the similarity. For instance, in recommender systems with implicit feedback, by treating unobserved pairs as negative, usually $\tilde{Y}_{ij} = 0$ or -1 is considered.

Clearly (5) involves a summation of mn terms, so the cost of directly applying an optimization algorithm is prohibitively proportional to $O(mn)$.

2.2 Linear Embedding Models

As mentioned in Section 1, most existing works of extreme similarity learning focus on linear embedding models. To avoid any $O(mn)$ cost occurs in solving the optimization problem, they consider a special kind of loss functions

$$\ell_{ij}(Y_{ij}, \hat{Y}_{ij}) = \begin{cases} \ell(Y_{ij}, \hat{Y}_{ij}) & (i, j) \in \mathbb{O}, \\ \frac{1}{2}\omega a_i b_j (\tilde{Y}_{ij} - \hat{Y}_{ij})^2 & (i, j) \notin \mathbb{O}, \end{cases} \quad (7)$$

where $\ell(a, b)$ is any non-negative and twice-differentiable function convex in b , and ω is a parameter for balancing the two kinds of losses. Two vectors $\mathbf{a} \in \mathbb{R}^m$, and $\mathbf{b} \in \mathbb{R}^n$ are chosen so that $a_i b_j$ is a cost associated with an unobserved pair. Besides, for the imputed label \tilde{Y}_{ij} , it is required [13, 28] that

$$\tilde{Y}_{ij} = \tilde{\mathbf{p}}_i^\top \tilde{\mathbf{q}}_j, \quad (8)$$

where $\tilde{\mathbf{p}}_i, \tilde{\mathbf{q}}_j$ are vectors output by the prior imputation model, and are fixed throughout training. For easy analysis, we assume that $\tilde{\mathbf{p}}_i, \tilde{\mathbf{q}}_j \in \mathbb{R}^k$ have the same length as $\mathbf{p}_i, \mathbf{q}_j$. With the loss in (7), past works were able to reduce the $O(mn)$ cost to $O(m+n)$. The main reason is because derived computations on all i and all j can be obtained by using values solely related to i and j , respectively. A conceptual illustration is in the following equation.

$$\sum_{i=1}^m \sum_{j=1}^n (\dots) = \sum_{i=1}^m (\dots) \times \sum_{j=1}^n (\dots). \quad (9)$$

The mainline of most past works is to incorporate more linear models into extreme similarity learning, which can range from simple matrix factorization [9, 16, 25] to complex models, e.g., matrix factorization with side information and (field-aware) factorization machines [3, 26, 29].

3 INFORMATION NEEDED IN OPTIMIZATION ALGORITHMS

Most optimization methods involve a procedure that repeatedly calculates the function value, the gradient, and/or other information. For example, a gradient descent method must calculate the negative gradient as the direction for an update. Crucially, from past developments for linear embeddings, the $O(mn)$ cost mainly occurs in these building blocks of optimization algorithms. The aim of this section is to derive $O(m+n)$ computation for these important components when nonlinear embeddings are applied.

3.1 From Linear to Nonlinear Embedding Models

We begin by discussing the differences between using linear and nonlinear embeddings. If linear embeddings are used, a key feature is that (4) becomes a multi-block convex problem. That is, (4) is reduced to a convex sub-problem when embedding variables of one side are fixed. Then existing works apply alternating least squares [16], (block) coordinate descent [3, 9, 25, 29], and alternating Newton method [26, 28] to sequentially solve each sub-problem by convex optimization techniques, which often iteratively calculate the following information.

- Function value,
- Gradient, and/or
- Hessian-vector product

Thus a focus was on deriving $O(m+n)$ operations for them.

If nonlinear embeddings are considered, in general (4) is not a block-convex problem. Thus a block-wise setting may not possess advantages so that optimization algorithms updating all variables together is a more natural choice. In any case, if we can obtain formulations for calculating the function value or the gradient over all variables, then the information over a block of variables is often in a simplified form. For example, the gradient over the block θ^u is a sub-vector of the whole gradient over $\theta = \begin{bmatrix} \theta^u \\ \theta^v \end{bmatrix}$. Therefore, in the rest of this section, we check the general scenario of considering all variables.

3.2 Evaluation of the Objective Value

The infeasible $O(mn)$ cost occurs if we directly compute $L(\theta)$ in (5). To handle this, our idea is to follow past works of linear embeddings to consider (7) as the loss function.

By applying (7), (5) is equivalent to

$$\begin{aligned} L(\theta) &= \sum_{(i,j) \in \mathbb{O}} \ell(Y_{ij}, \hat{Y}_{ij}) + \sum_{(i,j) \notin \mathbb{O}} \frac{1}{2}\omega a_i b_j (\tilde{Y}_{ij} - \hat{Y}_{ij})^2 + \lambda \Psi(\theta) \\ &= \underbrace{\sum_{(i,j) \in \mathbb{O}} \ell_{ij}^+(Y_{ij}, \hat{Y}_{ij})}_{L^+(\theta)} + \omega \underbrace{\sum_{i=1}^m \sum_{j=1}^n \ell_{ij}^-(Y_{ij}, \hat{Y}_{ij})}_{L^-(\theta)} + \lambda \Psi(\theta), \end{aligned} \quad (10)$$

where

$$\begin{aligned} \ell_{ij}^+(Y_{ij}, \hat{Y}_{ij}) &= \ell(Y_{ij}, \hat{Y}_{ij}) - \frac{1}{2}\omega a_i b_j (\tilde{Y}_{ij} - \hat{Y}_{ij})^2, \\ \ell_{ij}^-(Y_{ij}, \hat{Y}_{ij}) &= \frac{1}{2}a_i b_j (\tilde{Y}_{ij} - \hat{Y}_{ij})^2. \end{aligned} \quad (11)$$

The first term $L^+(\theta)$ involves the summation of $|\mathbb{O}|$ values, so with $|\mathbb{O}| \ll mn$ the bottleneck is on the second term $L^-(\theta)$.

Let $F(f)$ and $F(g)$ respectively represent the cost of operations related to the two nonlinear embedding models,¹ and

$$P = \begin{bmatrix} \mathbf{p}_1^\top \\ \vdots \\ \mathbf{p}_m^\top \end{bmatrix}, Q = \begin{bmatrix} \mathbf{q}_1^\top \\ \vdots \\ \mathbf{q}_n^\top \end{bmatrix}, \tilde{P} = \begin{bmatrix} \tilde{\mathbf{p}}_1^\top \\ \vdots \\ \tilde{\mathbf{p}}_m^\top \end{bmatrix}, \tilde{Q} = \begin{bmatrix} \tilde{\mathbf{q}}_1^\top \\ \vdots \\ \tilde{\mathbf{q}}_n^\top \end{bmatrix}, \quad (12)$$

where \tilde{P} and \tilde{Q} are constant matrices, but P and Q depend on θ . To compute $L^+(\theta)$ in (10), we first compute P and Q in $O(mF(f) + nF(g))$ time and store them in $O((m+n)k)$ space. Then for each $(i, j) \in \mathbb{O}$, from (3) and (8), we compute $\hat{Y}_{ij} = \mathbf{p}_i^\top \mathbf{q}_j$ and $\tilde{Y}_{ij} = \tilde{\mathbf{p}}_i^\top \tilde{\mathbf{q}}_j$ in $O(k)$ time. As a result, we can compute the entire $L^+(\theta)$ in $O(mF(f) + nF(g) + |\mathbb{O}|k)$ time.

For $L^-(\theta)$, we follow the idea in (9) to derive

$$L^-(\theta) = \frac{1}{2} \langle \tilde{P}_c, \tilde{Q}_c \rangle_F - \langle \tilde{P}_c, \tilde{Q}_c \rangle_F + \frac{1}{2} \langle P_c, Q_c \rangle_F, \quad (13)$$

where details are left in Appendix A.1. Later for gradient calculation we show details as a demonstration of using (9).

¹If neural networks are considered, operations such as forward or backward computation are considered. See more discussion in Section 5

In (13), $\langle \cdot, \cdot \rangle_F$ is a Frobenius inner product between two matrices,

$$\begin{aligned} P_c &= P^\top AP, \hat{P}_c = \tilde{P}^\top AP, \tilde{P}_c = \tilde{P}^\top A\tilde{P}, \\ Q_c &= Q^\top BQ, \hat{Q}_c = \tilde{Q}^\top BQ, \text{ and } \tilde{Q}_c = \tilde{Q}^\top B\tilde{Q}, \end{aligned} \quad (14)$$

where $A = \text{diag}(\mathbf{a}), B = \text{diag}(\mathbf{b})$ are two diagonal matrices. As $P \in \mathbb{R}^{m \times k}$ and $Q \in \mathbb{R}^{n \times k}$ have been pre-computed and cached in memory during computing $L^+(\theta)$, the matrices in (14) can be computed in $O((m+n)k^2)$ time and cached in $O(k^2)$ space. Then the Frobenius inner products between these matrices cost only $O(k^2)$ time.

From complexities of $L^+(\theta)$ and $L^-(\theta)$, the overall cost of evaluating $L(\theta)$ is

$$O(|\mathbb{O}|k + (m+n)k^2 + mF(f) + nF(g)). \quad (15)$$

3.3 Computation of Gradient Information

For $L(\theta)$ defined in (5), the gradient is

$$\nabla L(\theta) = \sum_{i=1}^m \sum_{j=1}^n \frac{\partial \hat{Y}_{ij}^\top}{\partial \theta} \frac{\partial \ell_{ij}^\top}{\partial \hat{Y}_{ij}} + \lambda \nabla \Psi(\theta). \quad (16)$$

Analogous to (5), it is impractical to directly compute $\nabla L(\theta)$ with infeasible $O(mn)$ costs. For derivatives with respect to a vector, we let $\partial \hat{Y}_{ij} / \partial \theta$ be a row vector, and later for operations such as $\partial f^i / \partial \theta^u$, we let it be a $k \times D_u$ matrix. In all other situations, a vector such as $\nabla L(\theta)$ is a column vector.

From (7) and (10), $\nabla L(\theta)$ in (16) is equivalent to

$$\nabla L(\theta) = \underbrace{\sum_{(i,j) \in \mathbb{O}} \frac{\partial \hat{Y}_{ij}^\top}{\partial \theta} \frac{\partial \ell_{ij}^\top}{\partial \hat{Y}_{ij}}}_{\nabla L^+(\theta)} + \omega \underbrace{\sum_{i=1}^m \sum_{j=1}^n \frac{\partial \hat{Y}_{ij}^\top}{\partial \theta} \frac{\partial \ell_{ij}^-}{\partial \hat{Y}_{ij}}}_{\nabla L^-(\theta)} + \lambda \nabla \Psi(\theta).$$

To calculate $\nabla L^+(\theta)$, let $X \in \mathbb{R}^{m \times n}$ be a sparse matrix with

$$X_{ij} = \begin{cases} \frac{\partial \hat{Y}_{ij}^\top}{\partial \theta} & (i, j) \in \mathbb{O}, \\ 0 & \text{otherwise.} \end{cases} \quad (17)$$

From (2) and (3),

$$\frac{\partial \hat{Y}_{ij}}{\partial \theta} = \left[\frac{\partial \hat{Y}_{ij}}{\partial f^i} \frac{\partial f^i}{\partial \theta^u}, \frac{\partial \hat{Y}_{ij}}{\partial g^j} \frac{\partial g^j}{\partial \theta^v} \right] = \left[\mathbf{q}_j^\top \frac{\partial f^i}{\partial \theta^u}, \mathbf{p}_i^\top \frac{\partial g^j}{\partial \theta^v} \right], \quad (18)$$

so with (17) we have

$$\nabla L^+(\theta) = \sum_{i=1}^m \sum_{j=1}^n \left(\frac{\partial \hat{Y}_{ij}^\top}{\partial \theta} X_{ij} \right) = \left[\sum_{i=1}^m \sum_{j=1}^n \frac{\partial f^i}{\partial \theta^u} \mathbf{q}_j X_{ij} \right]. \quad (19)$$

Let $\mathbf{J}^u \in \mathbb{R}^{m \times D_u \times k}, \mathbf{J}^v \in \mathbb{R}^{n \times D_v \times k}$ be tensors with

$$\text{the } i\text{-th slice } \mathbf{J}_{i,:}^u = \frac{\partial f^i}{\partial \theta^u} \text{ and the } j\text{-th slice } \mathbf{J}_{j,:}^v = \frac{\partial g^j}{\partial \theta^v}, \quad (20)$$

respectively. The top part of (19) can be computed by

$$\begin{aligned} \sum_{i=1}^m \sum_{j=1}^n \frac{\partial f^i}{\partial \theta^u} \mathbf{q}_j X_{ij} &= \sum_{i=1}^m \frac{\partial f^i}{\partial \theta^u} \left(\sum_{j=1}^n X_{ij} \mathbf{q}_j^\top \right)^\top \\ &= \sum_{i=1}^m \frac{\partial f^i}{\partial \theta^u} \left(X_{i,:} \mathbf{Q} \right)^\top = \langle \mathbf{J}^u, X\mathbf{Q} \rangle, \end{aligned} \quad (21)$$

where (21) is from (12) and the following definition of the tensor-matrix inner product

$$\langle \mathbf{J}, \mathbf{M} \rangle = \sum_{i=1}^m \mathbf{J}_{i,:} \langle \mathbf{M}_{i,:} \rangle^\top. \quad (22)$$

As the computation of the bottom part of (19) is similar, we omit the details. Then (19) can be written as

$$\nabla L^+(\theta) = \left[\langle \mathbf{J}^u, X\mathbf{Q} \rangle \right], \quad (23)$$

where P and Q are first computed in $O(mF(f) + nF(g))$ time and stored in $O((m+n)k)$ space. Then we compute $X\mathbf{Q}$ and $X^\top P$ in $O(|\mathbb{O}|k)$ time and $O((m+n)k)$ space.

To calculate $\nabla L^-(\theta)$, from (3) and (11) we have

$$\frac{\partial \ell_{ij}^-}{\partial \hat{Y}_{ij}} = a_i b_j (\mathbf{p}_i^\top \mathbf{q}_j - \tilde{\mathbf{p}}_i^\top \tilde{\mathbf{q}}_j).$$

This and (18) imply that

$$\begin{aligned} \nabla L^-(\theta) &= \sum_{i=1}^m \sum_{j=1}^n \left[\frac{\partial f^i}{\partial \theta^u} \mathbf{q}_j \right] a_i b_j (\mathbf{p}_i^\top \mathbf{q}_j - \tilde{\mathbf{p}}_i^\top \tilde{\mathbf{q}}_j) \\ &= \left[\sum_{i=1}^m \frac{\partial f^i}{\partial \theta^u} \mathbf{q}_j a_i b_j (\mathbf{q}_j^\top \mathbf{p}_i - \tilde{\mathbf{q}}_j^\top \tilde{\mathbf{p}}_i) \right] \\ &= \left[\sum_{j=1}^n \frac{\partial g^j}{\partial \theta^v} \mathbf{p}_i a_i b_j (\mathbf{p}_i^\top \mathbf{q}_j - \tilde{\mathbf{p}}_i^\top \tilde{\mathbf{q}}_j) \right] \\ &= \left[\sum_{i=1}^m \frac{\partial f^i}{\partial \theta^u} \left((\sum_{j=1}^n b_j \mathbf{q}_j \mathbf{q}_j^\top) a_i \mathbf{p}_i - (\sum_{j=1}^n b_j \mathbf{q}_j \tilde{\mathbf{q}}_j^\top) a_i \tilde{\mathbf{p}}_i \right) \right] \\ &= \left[\sum_{j=1}^n \frac{\partial g^j}{\partial \theta^v} \left((\sum_{i=1}^m a_i \mathbf{p}_i \mathbf{p}_i^\top) b_j \mathbf{q}_j - (\sum_{i=1}^m a_i \tilde{\mathbf{p}}_i \tilde{\mathbf{p}}_i^\top) b_j \tilde{\mathbf{q}}_j \right) \right], \end{aligned} \quad (24)$$

where the last two equalities follow from the idea in (9). From (14),

$$\begin{aligned} P_c &= \sum_{i=1}^m a_i \mathbf{p}_i \mathbf{p}_i^\top, Q_c = \sum_{j=1}^n b_j \mathbf{q}_j \mathbf{q}_j^\top, \\ \hat{P}_c &= \sum_{i=1}^m a_i \tilde{\mathbf{p}}_i \tilde{\mathbf{p}}_i^\top, \text{ and } \hat{Q}_c = \sum_{j=1}^n b_j \tilde{\mathbf{q}}_j \tilde{\mathbf{q}}_j^\top, \end{aligned} \quad (25)$$

so we can compute $\nabla L^-(\theta)$ by

$$\begin{aligned} \nabla L^-(\theta) &= \left[\sum_{i=1}^m \frac{\partial f^i}{\partial \theta^u} \mathbf{q}_j a_i (\mathbf{p}_i - \hat{Q}_c^\top a_i \tilde{\mathbf{p}}_i) \right] \\ &= \left[\sum_{j=1}^n \frac{\partial g^j}{\partial \theta^v} (\mathbf{P}_c b_j \mathbf{q}_j - \hat{P}_c^\top b_j \tilde{\mathbf{q}}_j) \right] \\ &= \left[\langle \mathbf{J}^u, AP_c Q_c - A\hat{P}_c \hat{Q}_c \rangle \right] \\ &= \left[\langle \mathbf{J}^v, BQ_c P_c - B\hat{Q}_c \hat{P}_c \rangle \right], \end{aligned} \quad (26)$$

where (26) follows from (12), (20), and (22). As P and Q have been pre-computed and cached in memory during computing $\nabla L^+(\theta)$, all matrices in (14) can be computed in $O((m+n)k^2)$ time and cached in $O(k^2)$ space. Then the right hand sides of $\langle \cdot, \cdot \rangle$ in (26) can be computed in $O((m+n)k^2)$ time as A and B are diagonal.

By combining (23) and (26), the entire $\nabla L(\theta)$ can be efficiently computed by

$$\nabla L(\theta) = \left[\langle \mathbf{J}^u, X\mathbf{Q} + \omega(AP_c Q_c - A\hat{P}_c \hat{Q}_c) \rangle \right] + \lambda \nabla \Psi(\theta), \quad (27)$$

which has a time complexity of

$$O(|\mathbb{O}|k + (m+n)k^2 + mF(f) + nF(g)). \quad (28)$$

The $mF(f) + nF(g)$ terms come from calculating P and Q in (23) and from $\partial f^i / \partial \theta^u, \forall i, \partial g^j / \partial \theta^v, \forall j$ in (20). The procedure of evaluating $\nabla L(\theta)$ is summarized in Algorithm 1.

From (15) and (28), the cost of one function evaluation is similar to that of one gradient evaluation.

Algorithm 1: Gradient evaluation: $\nabla L(\theta)$

Input: $\theta, \mathbb{O}, \tilde{P}, \tilde{Q}, A, B$.1 Calculate $P, Q, P_c, Q_c, \hat{P}_c, \hat{Q}_c$.²2 $X_{ij} \leftarrow \frac{\partial \ell_{ij}^+}{\partial \hat{Y}_{ij}}, \forall (i, j) \in \mathbb{O}$ 3 $X_q \leftarrow XQ, X_p \leftarrow X^\top P$ 4 $\nabla L(\theta) \leftarrow \left[\begin{array}{l} \langle \mathbf{J}^u, X_q + \omega(APQ_c - A\tilde{P}\tilde{Q}_c) \rangle \\ \langle \mathbf{J}^v, X_p + \omega(BQP_c - B\tilde{Q}\tilde{P}_c) \rangle \end{array} \right] + \lambda \nabla \Psi(\theta)$.**Output:** $\nabla L(\theta)$

3.4 Computation of Gauss-Newton Matrix-vector Products

For optimization methods using second-order information, commonly the product between Hessian and some vector \mathbf{d} is needed:

$$\nabla^2 L(\theta) \mathbf{d}. \quad (29)$$

The Hessian of (5) is

$$\nabla^2 L(\theta) = \sum_{i=1}^m \sum_{j=1}^n \frac{\partial \hat{Y}_{ij}}{\partial \theta} \frac{\partial^2 \ell_{ij}}{\partial \hat{Y}_{ij}^2} \frac{\partial \hat{Y}_{ij}}{\partial \theta} + \sum_{i=1}^m \sum_{j=1}^n \frac{\partial^2 \hat{Y}_{ij}}{\partial \theta^2} \frac{\partial \ell_{ij}}{\partial \hat{Y}_{ij}} + \lambda \nabla^2 \Psi(\theta). \quad (30)$$

Clearly, a direct calculation costs $O(mn)$. However, before reducing the $O(mn)$ cost to $O(m+n)$, we must address an issue that $\nabla^2 L(\theta)$ is not positive definite. In past studies of linear embeddings (e.g., [4, 26, 29]) the block-convex $L(\theta)$ implies that for a strictly convex sub-problem, the Hessian is positive definite. Without such properties, here we need a positive definite approximation of $\nabla^2 L(\theta)$. Following [15, 20], we remove the second term in (30) to have the following Gauss-Newton matrix [19]

$$G = \sum_{i=1}^m \sum_{j=1}^n \frac{\partial \hat{Y}_{ij}}{\partial \theta} \frac{\partial^2 \ell_{ij}}{\partial \hat{Y}_{ij}^2} \frac{\partial \hat{Y}_{ij}}{\partial \theta} + \lambda \nabla^2 \Psi(\theta), \quad (31)$$

which is positive definite from the convexity of ℓ_{ij} in \hat{Y}_{ij} and the strong convexity of $\Psi(\theta)$. The matrix-vector product becomes

$$G\mathbf{d} = \sum_{i=1}^m \sum_{j=1}^n \frac{\partial \hat{Y}_{ij}}{\partial \theta} \frac{\partial^2 \ell_{ij}}{\partial \hat{Y}_{ij}^2} \frac{\partial \hat{Y}_{ij}}{\partial \theta} \mathbf{d} + \lambda \nabla^2 \Psi(\theta) \mathbf{d}. \quad (32)$$

From (11) we have

$$\frac{\partial^2 \ell_{ij}^-}{\partial \hat{Y}_{ij}^2} = a_i b_j. \quad (33)$$

Following (10), we use (33) to re-write $G\mathbf{d}$ in (32) as

$$\underbrace{\sum_{(i,j) \in \mathbb{O}} \frac{\partial \hat{Y}_{ij}}{\partial \theta} \frac{\partial^2 \ell_{ij}^+}{\partial \hat{Y}_{ij}^2} \frac{\partial \hat{Y}_{ij}}{\partial \theta} \mathbf{d}}_{G^+ \mathbf{d}} + \omega \underbrace{\sum_{i=1}^m \sum_{j=1}^n a_i b_j \frac{\partial \hat{Y}_{ij}}{\partial \theta} \frac{\partial \hat{Y}_{ij}}{\partial \theta} \mathbf{d}}_{G^- \mathbf{d}} + \lambda \nabla^2 \Psi(\theta) \mathbf{d}. \quad (34)$$

Let $\mathbf{d} = \begin{bmatrix} \mathbf{d}^u \\ \mathbf{d}^v \end{bmatrix}$ be some vector. By defining

$$\mathbf{w}_i = \frac{\partial f^i}{\partial \theta^u} \mathbf{d}^u \in \mathbb{R}^k, \mathbf{h}_j = \frac{\partial g^j}{\partial \theta^v} \mathbf{d}^v \in \mathbb{R}^k, \quad (35)$$

²If for the same θ , $L(\theta)$ is calculated before $\nabla L(\theta)$, then these matrices have been stored and are available as input.

from (18), we have

$$\frac{\partial \hat{Y}_{ij}}{\partial \theta} \mathbf{d} = \mathbf{q}_j^\top \mathbf{w}_i + \mathbf{p}_i^\top \mathbf{h}_j. \quad (36)$$

To compute $G^+ \mathbf{d}$ in (34), we first compute

$$\mathbf{W} = \begin{bmatrix} \mathbf{w}_1^\top \\ \vdots \\ \mathbf{w}_m^\top \end{bmatrix} \in \mathbb{R}^{m \times k} \text{ and } \mathbf{H} = \begin{bmatrix} \mathbf{h}_1^\top \\ \vdots \\ \mathbf{h}_n^\top \end{bmatrix} \in \mathbb{R}^{n \times k}$$

in $O(mF(f) + nF(g))$ time and store them in $O((m+n)k)$ space. From (36), let $\mathbf{Z} \in \mathbb{R}^{m \times n}$ be a sparse matrix with

$$Z_{ij} = \begin{cases} \frac{\partial^2 \ell_{ij}^+}{\partial \hat{Y}_{ij}^2} \frac{\partial \hat{Y}_{ij}}{\partial \theta} \mathbf{d} = \frac{\partial^2 \ell_{ij}^+}{\partial \hat{Y}_{ij}^2} (\mathbf{q}_j^\top \mathbf{w}_i + \mathbf{p}_i^\top \mathbf{h}_j) & (i, j) \in \mathbb{O}, \\ 0 & \text{otherwise,} \end{cases}$$

which can be constructed in $O(|\mathbb{O}|k)$ time. Then from (19) and similar to the situation of computing (21) and (23), we have

$$G^+ \mathbf{d} = \sum_{(i,j) \in \mathbb{O}} \frac{\partial \hat{Y}_{ij}}{\partial \theta} \frac{\partial^2 \ell_{ij}^+}{\partial \hat{Y}_{ij}^2} \frac{\partial \hat{Y}_{ij}}{\partial \theta} \mathbf{d} = \sum_{i=1}^m \sum_{j=1}^n \frac{\partial \hat{Y}_{ij}}{\partial \theta} \mathbf{Z}_{ij} = \begin{bmatrix} \langle \mathbf{J}^u, \mathbf{ZQ} \rangle \\ \langle \mathbf{J}^v, \mathbf{Z}^\top \mathbf{P} \rangle \end{bmatrix}, \quad (37)$$

where \mathbf{ZQ} and $\mathbf{Z}^\top \mathbf{P}$ can be computed in $O(|\mathbb{O}|k)$ time and $O((m+n)k)$ space.

By using (18) and (36), $G^- \mathbf{d}$ in (34) can be computed by

$$\begin{aligned} G^- \mathbf{d} &= \sum_{i=1}^m \sum_{j=1}^n a_i b_j \frac{\partial \hat{Y}_{ij}}{\partial \theta} \frac{\partial \hat{Y}_{ij}}{\partial \theta} \mathbf{d} \\ &= \sum_{i=1}^m \sum_{j=1}^n \begin{bmatrix} \frac{\partial f^i}{\partial \theta^u} \mathbf{q}_j \\ \frac{\partial g^j}{\partial \theta^v} \mathbf{p}_i \end{bmatrix} a_i b_j (\mathbf{w}_i^\top \mathbf{q}_j + \mathbf{p}_i^\top \mathbf{h}_j) \quad (38) \\ &= \begin{bmatrix} \langle \mathbf{J}^u, \mathbf{AWQ}_c + \mathbf{APH}_c \rangle \\ \langle \mathbf{J}^v, \mathbf{BHP}_c + \mathbf{BQW}_c \rangle \end{bmatrix}, \quad (39) \end{aligned}$$

where (39) is by a similar derivation from (24) to (26), and by first computing the following $k \times k$ matrices in $O((m+n)k^2)$ time:

$$\mathbf{W}_c = \sum_{i=1}^m a_i \mathbf{w}_i \mathbf{p}_i^\top = \mathbf{W}^\top \mathbf{A} \mathbf{P}, \quad \mathbf{H}_c = \sum_{j=1}^n b_j \mathbf{h}_j \mathbf{q}_j^\top = \mathbf{H}^\top \mathbf{B} \mathbf{Q}.$$

Combining (37) and (39), we can compute (34) by

$$G\mathbf{d} = \begin{bmatrix} \langle \mathbf{J}^u, \mathbf{ZQ} + \omega(\mathbf{AWQ}_c + \mathbf{APH}_c) \rangle \\ \langle \mathbf{J}^v, \mathbf{Z}^\top \mathbf{P} + \omega(\mathbf{BHP}_c + \mathbf{BQW}_c) \rangle \end{bmatrix} + \lambda \nabla^2 \Psi(\theta) \mathbf{d}. \quad (40)$$

Thus the total time complexity of (40) is

$$O(|\mathbb{O}|k + (m+n)k^2 + 2mF(f) + 2nF(g)),$$

where the term ‘2’ comes from the operations in (35) and (38). More details are in Section 5.

4 OPTIMIZATION METHODS FOR EXTREME SIMILARITY LEARNING WITH NONLINEAR EMBEDDINGS

With the efficient computation of key components given in Section 3, we demonstrate that many optimization algorithms can be used. While block-wise minimization is applicable, we mentioned in Section 3.1 that without a block-convex $L(\theta)$, such a setting may not be advantageous. Thus in our discussion we focus more on standard optimization techniques that update all variables at a time.

4.1 (Stochastic) Gradient Descent Method

The standard gradient descent method considers the negative gradient direction

$$\mathbf{s} = -\nabla L(\boldsymbol{\theta}) \quad (41)$$

and update $\boldsymbol{\theta}$ by

$$\boldsymbol{\theta} \leftarrow \boldsymbol{\theta} + \delta \mathbf{s}, \quad (42)$$

where δ is the step size. To ensure the convergence, usually a line-search procedure is conducted to find δ satisfying the following sufficient decrease condition

$$L(\boldsymbol{\theta} + \delta \mathbf{s}) \leq L(\boldsymbol{\theta}) + \eta \delta \mathbf{s}^\top \nabla L(\boldsymbol{\theta}), \quad (43)$$

where $\eta \in (0, 1)$ is a pre-defined constant. During line search, given any $\boldsymbol{\theta} + \delta \mathbf{s}$, the corresponding function value must be calculated.

Combined with the cost of $\nabla L(\boldsymbol{\theta})$ in (28) and the cost of $L(\boldsymbol{\theta})$ in (15), the complexity of each gradient descent iteration is

$$O((\#LS + 1) \times (|\mathbb{O}|k + (m+n)k^2 + mF(f) + nF(g))), \quad (44)$$

where the term ‘1’ in $(\#LS + 1)$ comes from the cost of $\nabla L(\boldsymbol{\theta})$, and $\#LS$ is the number of line search steps.

In some machine learning applications, the gradient calculation is considered too expensive so an approximation using partial data samples is preferred. This leads to the stochastic gradient (SG) method. Without the full gradient, the update rule (42) may no longer lead to the decrease of the function value. Thus a line search procedure is not useful and the step size δ is often decided by some pre-defined rules or tuned as a hyper-parameter.

Unfortunately, it is known that for extreme similarity learning, a direct implementation of SG methods faces the difficulty of $O(mn)$ cost (e.g., [25]). Now all (i, j) pairs are considered as our data set and the cost of each SG step is proportional to the number of sampled pairs. Thus to go over all or most pairs many SG steps are needed and the cost is $O(mn)$. In contrast, by techniques in Section 3.3 to calculate the gradient in $O(m+n)$ cost, all pairs have been considered. A natural remedy is to incorporate techniques in Section 3.3 to the SG method. Instead of randomly selecting pairs in an SG step, [13] proposed to select pairs in a set $\mathbb{B} = \hat{\mathbb{U}} \times \hat{\mathbb{V}}$ where $\hat{\mathbb{U}} \subseteq \mathbb{U}$ and $\hat{\mathbb{V}} \subseteq \mathbb{V}$. If $\hat{m} = |\hat{\mathbb{U}}|$ and $\hat{n} = |\hat{\mathbb{V}}|$, then by slightly modifying the derivation in (27), calculating the subsampled gradient involves $O(\hat{m} + \hat{n})$ instead of $O(m\hat{n})$ cost. Analogous to (28), the subsampled gradient on \mathbb{B} can be calculated in

$$O(|\hat{\mathbb{O}}|k + (\hat{m} + \hat{n})k^2 + \hat{m}F(f) + \hat{n}F(g)),$$

where $\hat{\mathbb{O}}$ is the subset of observed pairs falling in $\hat{\mathbb{U}} \times \hat{\mathbb{V}}$. Then the complexity of each data pass of handling mn pairs via $\frac{mn}{\hat{m}\hat{n}}$ SG steps is

$$O(|\mathbb{O}|k + (\frac{mn}{\hat{n}} + \frac{mn}{\hat{m}})k^2 + \frac{mn}{\hat{n}}F(f) + \frac{mn}{\hat{m}}F(g)). \quad (45)$$

Thus the $O(mn)$ cost is deducted by a factor.

We consider another SG method that has been roughly discussed in [13]. The idea is to note that if $\mathbb{O}_{i,:} = \{j : (i, j) \in \mathbb{O}\}$ and $\mathbb{O}_{:,j} = \{i : (i, j) \in \mathbb{O}\}$, then

$$\begin{aligned} \sum_{i=1}^m \sum_{j=1}^n a_i b_j (\dots) &= \sum_{i=1}^m \sum_{j' \in \mathbb{O}_{i,:}} \sum_{j=1}^n \sum_{i' \in \mathbb{O}_{:,j}} \frac{a_i}{|\mathbb{O}_{i,:}|} \frac{b_j}{|\mathbb{O}_{:,j}|} (\dots) \\ &= \sum_{(i,j') \in \mathbb{O}} \sum_{(i',j) \in \mathbb{O}} \frac{a_i}{|\mathbb{O}_{i,:}|} \frac{b_j}{|\mathbb{O}_{:,j}|} (\dots). \end{aligned} \quad (46)$$

Then instead of selecting a subset from mn pairs, the setting becomes to select two independent sets $\hat{\mathbb{O}}_1, \hat{\mathbb{O}}_2 \subset \mathbb{O}$ for constructing the subsampled gradient, where $\hat{\mathbb{O}}_1$ and $\hat{\mathbb{O}}_2$ respectively corresponding to the first and the second summations in (46). Therefore, the task of going over mn pairs is replaced by selecting $\hat{\mathbb{O}}_1 \subset \mathbb{O}$, and $\hat{\mathbb{O}}_2 \subset \mathbb{O}$ at a time to cover the whole $\mathbb{O} \times \mathbb{O}$ after many SG steps. We leave details of deriving the subsampled gradient in Appendix A.2, and the complexity for the task comparable to going over all mn pairs by standard SG methods is

$$O\left(\left(\frac{|\mathbb{O}|^2}{|\hat{\mathbb{O}}_1|} + \frac{|\mathbb{O}|^2}{|\hat{\mathbb{O}}_2|}\right)(k + k^2 + F(f) + F(g))\right). \quad (47)$$

A comparison with (45) shows how (47) might alleviate the $O(mn)$ cost. This relies on $|\mathbb{O}| \ll mn$ and some appropriate choices of $|\hat{\mathbb{O}}_1|$ and $|\hat{\mathbb{O}}_2|$. We refer to this method as the SG method based on (46).

The above setting is further extended in [13] to Stochastic Online Gramian (SOGram), which applies a variance reduction scheme to improve the subsampled gradient. Details are in Appendix A.3.

We compare the above methods in Table 2 by checking the cost per data pass, which roughly indicates the process of handling mn pairs. Note that the total cost also depends on the number of data passes. In experiments we will see the trade-off between the cost per data pass and the number of total passes.

4.2 Methods Beyond Gradient Descent

We can further consider methods that incorporate the second-order information. By considering the following second-order Taylor expansion

$$L(\boldsymbol{\theta} + \mathbf{s}) \approx L(\boldsymbol{\theta}) + \nabla L(\boldsymbol{\theta})^\top \mathbf{s} + \frac{1}{2} \mathbf{s}^\top \nabla^2 L(\boldsymbol{\theta}) \mathbf{s}, \quad (48)$$

Newton methods find a search direction \mathbf{s} by minimizing (48). However, this way of finding \mathbf{s} is often complicated and expensive. In the rest of this section, we discuss some practically viable options to use the second-order information.

4.2.1 Gauss-Newton Method. It is mentioned in Section 3.4 that Gauss-Newton matrix G in (31) is a positive-definite approximation of $\nabla^2 L(\boldsymbol{\theta})$. Thus replacing $\nabla^2 L(\boldsymbol{\theta})$ in (48) with G leads to a convex quadratic optimization problem with the solution satisfying

$$G\mathbf{s} = -\nabla L(\boldsymbol{\theta}). \quad (49)$$

However, for large-scale applications, not only is G in (31) too large to be stored in memory, but the matrix inversion is also computationally expensive.

Following [14, 15], we may solve (49) by the conjugate gradient (CG) method [10], which involves a sequence of matrix-vector products $G\mathbf{d}$ between G and some vector \mathbf{d} . The derivation in Section 3.4 has shown how to calculate $G\mathbf{d}$ without explicitly forming and storing G . Therefore, Gauss-Newton methods are a feasible approach for extreme similarity learning without the $O(mn)$ cost.

After a Newton direction \mathbf{s} is obtained, the line search procedure discussed in Section 4.1 is conducted to find a suitable δ to ensure the convergence. The complexity of each Newton iteration is

$$\begin{aligned} O(\#CG \times (|\mathbb{O}|k + (m+n)k^2 + 2mF(f) + 2nF(g)) \\ + (\#LS + 1) \times (|\mathbb{O}|k + (m+n)k^2 + mF(f) + nF(g))), \end{aligned} \quad (50)$$

Table 2: A comparison of gradient-based methods on time complexity of each data pass. For the last setting, $|\hat{\mathcal{O}}| = \min(|\hat{\mathcal{O}}_1|, |\hat{\mathcal{O}}_2|)$.

Method	Time complexity per data pass
Naive SG	$mnF(f) + mnF(g)$
Gradient (Alg. 1)	$ \mathcal{O} k + (m+n)k^2 + mF(f) + nF(g)$
Selecting $\mathbb{B} = \hat{\mathbf{U}} \times \hat{\mathbf{V}}$	$ \mathcal{O} k + (\frac{n}{n}m + \frac{m}{m}n)k^2 + \frac{n}{n}mF(f) + \frac{m}{m}nF(g)$
SG based on (46)	$\frac{ \mathcal{O} ^2}{ \hat{\mathcal{O}} }k + \frac{ \mathcal{O} ^2}{ \hat{\mathcal{O}} }k^2 + \frac{ \mathcal{O} ^2}{ \hat{\mathcal{O}} }F(f) + \frac{ \mathcal{O} ^2}{ \hat{\mathcal{O}} }F(g)$

where #CG is numbers of CG steps. The overall procedures of the CG method and the Gauss-Newton method are summarized in Appendix A.4.

4.2.2 Diagonal Scaling Methods. From (50), the CG procedure may be expensive so a strategy is to replace G in (49) with a positive-definite diagonal matrix, which becomes a very loose approximation of $\nabla^2 L(\theta)$. In this case, s can be efficiently obtained by dividing each value in $-\nabla L(\theta)$ by the corresponding diagonal element. Methods following this idea are referred to as diagonal scaling methods.

For example, the AdaGrad algorithm [6] is a popular diagonal scaling method. Given $\nabla L(\theta)$ of current θ , AdaGrad maintains a zero-initialized diagonal matrix $M \in \mathbb{R}^{D \times D}$ by

$$M \leftarrow M + \text{diag}(\nabla L(\theta))^2.$$

Then by considering $(\mu I + M)^{\frac{1}{2}}$, where I is the identity matrix and $\mu > 0$ is a small constant to ensure the invertibility, the direction is

$$s = -(\mu I + M)^{-\frac{1}{2}} \nabla L(\theta). \quad (51)$$

In comparison with the gradient descent method, the extra cost is minor in $\mathcal{O}(D)$, where $D = D_u + D_v$ is the total number of variables. Instead of using the full $\nabla L(\theta)$, this method can be extended to use subsampled gradient by the SG framework discussed in Section 4.1.

5 IMPLEMENTATION DETAILS

So far we used $F(f)$ and $F(g)$ to represent the cost of operations on nonlinear embeddings. Here we discuss more details if automatic differentiation is employed for easily trying different network architectures. To begin, note that $[\partial f^i / \partial \theta^u, \partial g^j / \partial \theta^v]$ is the Jacobian of $(f^i(\theta^u), g^j(\theta^v))$. Transposed Jacobian-vector products are used in $\nabla L(\theta)$ for all methods and Gd for the Gauss-Newton method. In fact, (27) and (40) are almost in the same form. It is known that for such an operation, the reverse-mode automatic differentiation (i.e., back-propagation) should be used so that only one pass of the network is needed [2]. For feedforward networks, including extensions such as convolutional networks, the cost of a backward process is within a constant factor of the forward process for computing the objective value (e.g., Section 5.2 of [21]). The remaining major operation to be discussed is the Jacobian-vector product in (35) in Section 4.2.1. This operation can be implemented with the forward-mode automatic differentiation in just one pass [2], and the cost is also within a constant factor of the cost of forward objective evaluation.³ We conclude that if neural networks are used, using $F(f)$ and $F(g)$ as the cost of nonlinear embeddings is suitable.

³See, for example, the discussion at https://www.csie.ntu.edu.tw/~cjlin/courses/optdl2021/slides/newton_gauss_newton.pdf.

Table 3: Data statistics and the selected hyper-parameters.

Data set	m	n	$ \mathcal{O} $	$\log_2 \omega$	$\log_2 \lambda$
m11m	6,037	3,513	517,770	-4	2
m110m	69,786	10,210	4,505,820	-8	-2
netflix	478,251	17,768	51,228,351	-8	0
wiki-simple	85,270	55,695	5,478,549	-10	2

Neural networks are routinely trained on GPU instead of CPU for better efficiency, but the smaller memory of GPU causes difficulties for larger problems or network architectures. This situation is even more serious for our case because some other matrices besides the two networks are involved; see (27) and (40). We develop some techniques to make large-scale training feasible, where details are left in supplementary materials.

6 EXPERIMENTS

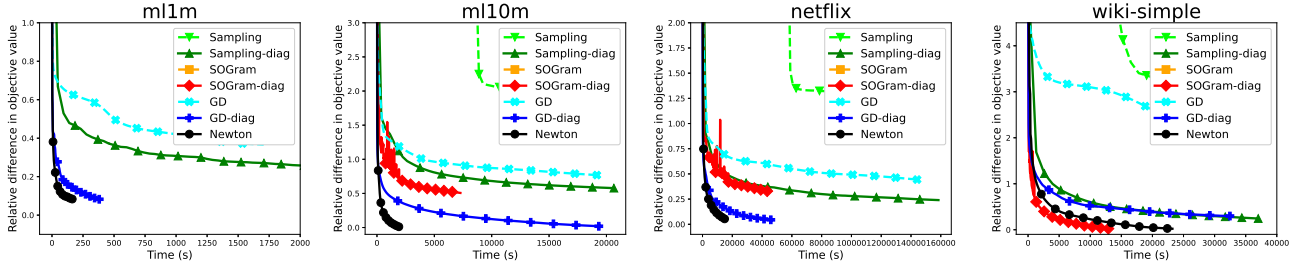
In this section, after presenting our experimental settings, we compare several methods discussed in Section 4 in terms of the convergence speed on objective value decreasing and achieving the final test performance.

6.1 Experimental Settings


6.1.1 Data sets. We consider three data sets of recommender systems with implicit feedback and one data set of link predictions. For data sets of recommender systems with implicit feedback, the target similarity of each observed user-item pair is 1, while those of the unobserved pairs are unrevealed. We consider m11m, m110m, and netflix provided by [25]. For the data set of link predictions, we follow [13] to consider the problem of learning the intra-site links between Wikipedia pages, where the target similarity of each observed page_{from}-page_{to} pair is 1 if there is a link from page_{from} to page_{to}, and unrevealed otherwise. However, because the date information of the dump used in [13] is not provided, their generated sets are irreproducible. We use the latest Wikipedia graph⁴ on pages in simple English and denote the generated data set as wiki-simple. For m11m, m110m and netflix, training and test sets are available. For wiki-simple, we follow [13] to use a 9-to-1 ratio for the training/test split. The statistics of all data sets are listed in Table 3. Note that the aim here is to check the convergence behavior, so we do not need a validation set for hyper-parameter selection.

6.1.2 Model and Hyper-parameters. We train a two-tower neural network as in Figure 1. Both towers have three fully-connected layers, where the first two layers contain 256 hidden units equipped with ELU [5] activation, and the last layer contains $k = 128$ hidden units without activation. For all experiments, we choose the logistic loss $\ell(y, \hat{y}) = \log(1 + \exp(-y\hat{y}))$, an L2 regularizer $\Psi(\theta) = \frac{1}{2} \|\theta\|_2^2$, and uniform weights $a_i = 1, \forall i = 1, \dots, m$, $b_j = 1, \forall j = 1, \dots, n$ in (7). For the imputed labels \tilde{Y} , we follow past works [25, 26] to apply a constant -1 for any unobserved pair by setting $\tilde{p}_i = -1/\sqrt{k}$, and $\tilde{q}_j = 1/\sqrt{k}$. For ω and λ , for each data set, we consider a grid of hyper-parameter combinations and select the one achieving the best results on the test set. The selected ω and λ for each data set are listed in Table 3.

⁴<https://dumps.wikimedia.org/simplewiki/20200901/>



(a) Relative difference in objective value versus training time.



(b) MAP@5 versus training time.

Figure 2: Comparison of different algorithms on four data sets. SOGram and Sampling may be too slow to be shown.

6.1.3 *Compared Optimization Methods.* We compare seven optimization methods for extreme similarity learning, which are categorized into two groups. The first group includes the following three gradient-based methods discussed in Section 4.1.

- GD: This is the gradient descent method. For each step, we compute $\mathbf{s} = -\nabla L(\boldsymbol{\theta})$ by (27). To ensure the convergence, backtracking line search is conducted to find the step size δ .
- SOGram: This is a stochastic gradient method proposed in [13]; see the discussion in Appendix A.3. We set $|\hat{\mathbf{O}}_1| = |\hat{\mathbf{O}}_2| = \rho|\mathbf{O}|$, where $\rho = 10^{-2}$ is the sampling ratio of observed pairs. We follow [13] to set $\alpha = 0.1$, which is a hyper-parameter of controlling the bias-variance tradeoff.
- Sampling: This is the stochastic gradient method of choosing $\mathbb{B} = \hat{\mathbf{U}} \times \hat{\mathbf{V}}$ [13]. To make the number of steps in each data pass comparable to SOGram, for Sampling we set $\hat{m} = \rho m$ and $\hat{n} = \rho n$, where ρ is the ratio used in SOGram.

The step size δ in SOGram and Sampling is set to 2^{-25} , 2^{-23} , 2^{-25} and 2^{-25} respectively for ml1m, ml10m, netflix and wiki-simple. These values are from GD’s most used δ .

The second group includes the following four methods incorporating the second-order information discussed in Section 4.2.

- Newton: This is the Gauss-Newton method proposed in Section 4.2.1. For CG, we set the stopping tolerance $\xi = 0.1$ in Appendix Algorithm 3 and 30 as the maximal number of CG steps.
- GD-dia, SOGram-dia, Sampling-dia: These three methods are respectively extended from GD, SOGram, and Sampling by applying the diagonal scaling in (51). For GD-dia, the same as GD, it finds δ by conducting backtracking line search. For SOGram-dia and Sampling-dia, we set $\delta = 0.01$.

For GD, Newton, and GD-dia, which conduct line search, we set $\eta = 10^{-4}$ for the sufficient decrease condition (43) and sequentially check δ in $\{\delta_{\text{init}}, \frac{\delta_{\text{init}}}{2}, \frac{\delta_{\text{init}}}{4}, \dots\}$, where δ_{init} is the step size taken in the last iteration, and is doubled every five iterations to avoid being stuck with a small step size.

6.1.4 *Environment and Implementation.* We use TensorFlow [1] compiled with Intel® Math Kernel Library (MKL) to implement all algorithms. Because most operations are matrix-based and MKL is optimized for such operations, our implementation should be sufficiently efficient. For sparse operations involving iterating \mathbf{O} (e.g., $L^+(\boldsymbol{\theta})$ and XQ), we implement them by NumPy [8] and a C extension, where we parallelize these operations by OpenMP. As our applied neural networks described in Section 6.1.2 are not complex, we empirically do not observe many advantages of GPU over CPU on the running time. Thus all experiments are conducted on a Linux machine with 10 cores of Intel® Core™ i7-6950X CPUs and 128GB memory.

6.1.5 *Evaluation Criterion of Test Performance.* On test sets, we report mean average precision (MAP) on top-5 ranked entities. For left entity i , let $\tilde{\mathbf{O}}_{i,:}$ be the set of its similar right entities in the test set, and $\tilde{\mathbf{O}}_{i,:K}^{\text{sorted}}$ be the set of top- K right entities with the highest predicted similarity. We define

$$\text{MAP@5} = \frac{1}{5} \sum_{K=1}^5 \frac{1}{\tilde{m}} \sum_{i=1}^{\tilde{m}} \frac{1}{K} |\tilde{\mathbf{O}}_{i,:} \cap \tilde{\mathbf{O}}_{i,:K}^{\text{sorted}}|, \quad (52)$$

where \tilde{m} is the number of left entities in the test set.

6.2 A Comparison on the Convergence Speed

We first investigate the relationship between the running time and the relative difference in objective value

$$(L(\theta) - L^*)/L^*,$$

where L^* is the lowest objective value reached among all settings. The result is presented in Figure 2a.

For the three gradient-based methods, SOGram, and Sampling are slower than GD. This situation is different from that of training general neural networks, where SG are the dominant methods. The reason is that though Sampling and SOGram take more iterations than GD in each data pass, from Table 2, they have much higher complexity per data pass.

For GD-diag, SOGram-diag, Sampling-diag, and Newton, which incorporate the second-order information, they are consistently faster than gradient-based methods. Specifically, between Newton and GD, the key difference is that Newton additionally considers the second-order information for yielding each direction. The great superiority of Newton confirms the importance of leveraging the second-order information. By comparing GD-diag and GD, we observe that GD-diag is much more efficient. Though both have similar complexities for obtaining search directions, GD-diag leverages partial second-order information from AdaGrad’s diagonal scaling. However, the effect of the diagonal scaling is limited, so GD-diag is slower than Newton. Finally, GD-diag is still faster than SOGram-diag, Sampling-diag as in the case of without diagonal scaling.

Next, we present in Figure 2b the relationship between the running time and MAP@5 evaluated on test sets. It is observed that a method with faster convergence in objective values is also faster in MAP@5.

7 CONCLUSIONS

In this work, we study extreme similarity learning with nonlinear embeddings. The goal is to alleviate the $O(mn)$ cost in the training process. While this topic has been well studied for the situation of using linear embeddings, a systematic study for nonlinear embeddings was lacking. We fill the gap in the following aspects. First, for important operations in optimization algorithms such as function and gradient evaluation, clean formulations with $O(m+n)$ cost are derived. Second, these formulations enable the use of many optimization algorithms for extreme similarity learning with nonlinear embedding models. We detailedly discuss some of them and check implementation issues. Experiments show that efficient training by some algorithms can be achieved.

ACKNOWLEDGMENTS

We would like to thank Hong Zhu, Yaxu Liu, and Jui-Nan Yen for helpful discussions. This work was supported by MOST of Taiwan grant 107-2221-E-002-167-MY3.

REFERENCES

- [1] Martin Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, Manjunath Kudlur, Josh Levenberg, Rajat Monga, Sherry Moore, Derek G. Murray, Benoit Steiner, Paul Tucker, Vijay Vasudevan, Pete Warden, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. 2016. TensorFlow: A system for large-scale machine learning. In *OSDI*.
- [2] Atilim Gunes Baydin, Barak A. Pearlmutter, Alexey Andreyevich Radul, and Jeffrey Mark Siskind. 2018. Automatic differentiation in machine learning: a survey. *JMLR* 18 (2018), 1–43.
- [3] Immanuel Bayer, Xiangnan He, Bhargav Kanagal, and Steffen Rendle. 2017. A generic coordinate descent framework for learning from implicit feedback. In *WWW*.
- [4] Wei-Sheng Chin, Bo-Wen Yuan, Meng-Yuan Yang, and Chih-Jen Lin. 2018. An Efficient Alternating Newton Method for Learning Factorization Machines. *ACM TIST* 9 (2018), 72:1–72:31.
- [5] Djork-Arné Clevert, Thomas Unterthiner, and Sepp Hochreiter. 2016. Fast and Accurate Deep Network Learning by Exponential Linear Units (ELUs). In *ICLR*.
- [6] John Duchi, Elad Hazan, and Yoram Singer. 2011. Adaptive Subgradient Methods for Online Learning and Stochastic Optimization. *JMLR* 12 (2011), 2121–2159.
- [7] Ali Mamdouh Elkahky, Yang Song, and Xiaodong He. 2015. A multi-view deep learning approach for cross domain user modeling in recommendation systems. In *WWW*.
- [8] Charles R. Harris, K. Jarrod Millman, Stéfán J. van der Walt, Ralf Gommers, Pauli Virtanen, David Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J Smith, et al. 2020. Array programming with NumPy. *Nature* 585, 7825 (2020), 357–362.
- [9] Xiangnan He, Hanwang Zhang, Min-Yen Kan, and Tat-Seng Chua. 2016. Fast matrix factorization for online recommendation with implicit feedback. In *SIGIR*.
- [10] Magnus Rudolph Hestenes and Eduard Stiefel. 1952. Methods of Conjugate Gradients for Solving Linear Systems. *JRNBS* 49 (1952), 409–436.
- [11] Jui-Ting Huang, Ashish Sharma, Shuying Sun, Li Xia, David Zhang, Philip Pronin, Janani Padmanabhan, Giuseppe Ottaviano, and Linjun Yang. 2020. Embedding-based Retrieval in Facebook Search. In *KDD*.
- [12] Po-Sen Huang, Xiaodong He, Jianfeng Gao, Li Deng, Alex Acero, and Larry Heck. 2013. Learning deep structured semantic models for web search using clickthrough data. In *CIKM*.
- [13] Walid Krichene, Nicolas Mayoraz, Steffen Rendle, Li Zhang, Xinyang Yi, Lichan Hong, Ed Chi, and John Anderson. 2019. Efficient Training on Very Large Corpora via Gramian Estimation. In *ICLR*.
- [14] Chih-Jen Lin, Ruby C. Weng, and S. Sathya Keerthi. 2008. Trust region Newton method for large-scale logistic regression. *JMLR* 9 (2008), 627–650.
- [15] James Martens. 2010. Deep learning via Hessian-free optimization. In *ICML*.
- [16] Rong Pan and Martin Scholz. 2009. Mind the Gaps: Weighting the Unknown in Large-scale One-class Collaborative Filtering. In *KDD*.
- [17] Rong Pan, Yunhong Zhou, Bin Cao, Nathan N Liu, Rajan Lukose, Martin Scholz, and Qiang Yang. 2008. One-class collaborative filtering. In *ICDM*.
- [18] Steffen Rendle, Walid Krichene, Li Zhang, and John Anderson. 2020. Neural Collaborative Filtering vs. Matrix Factorization Revisited. In *RecSys*.
- [19] Nicol N. Schraudolph. 2002. Fast curvature matrix-vector products for second-order gradient descent. *Neural Comput.* 14 (2002), 1723–1738.
- [20] Chien-Chih Wang, Chun-Heng Huang, and Chih-Jen Lin. 2015. Subsampled Hessian Newton Methods for Supervised Learning. *Neural Comput.* 27 (2015), 1766–1795.
- [21] Chien-Chih Wang, Kent Loong Tan, and Chih-Jen Lin. 2020. Newton Methods for Convolutional Neural Networks. *ACM TIST* 11 (2020), 19:1–19:30.
- [22] Xiaojie Wang, Rui Zhang, Yu Sun, and Jianzhong Qi. 2019. Doubly robust joint learning for recommendation on data missing not at random. In *ICML*.
- [23] Yinfei Yang, Steve Yuan, Daniel Cer, Sheng-Yi Kong, Noah Constant, Petr Pilar, Heming Ge, Yun-Hsuan Sung, Brian Strope, and Ray Kurzweil. 2018. Learning Semantic Textual Similarity from Conversations. In *ReplANLP*.
- [24] Xinyang Yi, Ji Yang, Lichan Hong, Derek Zhiyuan Cheng, Lukasz Heldt, Aditee Kumthekar, Zhe Zhao, Li Wei, and Ed Chi. 2019. Sampling-Bias-Corrected Neural Modeling for Large Corpus Item Recommendations. In *RecSys*.
- [25] Hsiang-Fu Yu, Mikhail Bilenko, and Chih-Jen Lin. 2017. Selection of Negative Samples for One-class Matrix Factorization. In *SDM*.
- [26] Hsiang-Fu Yu, Hsin-Yuan Huang, Inderjit S. Dhillon, and Chih-Jen Lin. 2017. A Unified Algorithm for One-class Structured Matrix Factorization with Side Information. In *AAAI*.
- [27] Hsiang-Fu Yu, Prateek Jain, Purushottam Kar, and Inderjit S. Dhillon. 2014. Large-scale Multi-label Learning with Missing Labels. In *ICML*.
- [28] Bowen Yuan, Jui-Yang Hsia, Meng-Yuan Yang, Hong Zhu, Chihyao Chang, Zhenhua Dong, and Chih-Jen Lin. 2019. Improving Ad Click Prediction by Considering Non-displayed Events. In *CIKM*.
- [29] Bowen Yuan, Meng-Yuan Yang, Jui-Yang Hsia, Hong Zhu, Zhirong Liu, Zhenhua Dong, and Chih-Jen Lin. 2019. *One-class Field-aware Factorization Machines for Recommender Systems with Implicit Feedbacks*. Technical Report. National Taiwan Univ.
- [30] Kai Zhong, Zhao Song, Prateek Jain, and Inderjit S Dhillon. 2019. Provable Non-linear Inductive Matrix Completion. In *NIPS*.

A APPENDIX

A.1 Detailed Derivation of $L^-(\theta)$

To compute $L^-(\theta)$, we rewrite it as

$$\begin{aligned}
L^-(\theta) &= \frac{1}{2} \sum_{i=1}^m \sum_{j=1}^n a_i b_j (\tilde{\mathbf{p}}_i^\top \tilde{\mathbf{q}}_j - \mathbf{p}_i^\top \mathbf{q}_j)^2 \\
&= \frac{1}{2} \sum_{i=1}^m \sum_{j=1}^n a_i b_j \left(\tilde{\mathbf{p}}_i^\top \tilde{\mathbf{q}}_j (\tilde{\mathbf{q}}_j^\top \tilde{\mathbf{p}}_i) - 2 \tilde{\mathbf{p}}_i^\top \tilde{\mathbf{q}}_j (\mathbf{q}_j^\top \mathbf{p}_i) + \mathbf{p}_i^\top \mathbf{q}_j (\mathbf{q}_j^\top \mathbf{p}_i) \right) \\
&= \frac{1}{2} \sum_{i=1}^m \left(a_i \tilde{\mathbf{p}}_i^\top \left(\sum_{j=1}^n b_j \tilde{\mathbf{q}}_j \tilde{\mathbf{q}}_j^\top \right) \tilde{\mathbf{p}}_i - 2 a_i \tilde{\mathbf{p}}_i^\top \left(\sum_{j=1}^n b_j \tilde{\mathbf{q}}_j \mathbf{q}_j^\top \right) \mathbf{p}_i \right. \\
&\quad \left. + a_i \mathbf{p}_i^\top \left(\sum_{j=1}^n b_j \mathbf{q}_j \mathbf{q}_j^\top \right) \mathbf{p}_i \right) \\
&= \frac{1}{2} \langle \tilde{\mathbf{P}}_c, \tilde{\mathbf{Q}}_c \rangle_F - \langle \hat{\mathbf{P}}_c, \hat{\mathbf{Q}}_c \rangle_F + \frac{1}{2} \langle \mathbf{P}_c, \mathbf{Q}_c \rangle_F. \tag{53}
\end{aligned}$$

To calculate (53), we need matrices listed in (14). For $\tilde{\mathbf{P}}_c$ and $\tilde{\mathbf{Q}}_c$, they are constant matrices throughout the training process, so we can calculate them and $\langle \tilde{\mathbf{P}}_c, \tilde{\mathbf{Q}}_c \rangle_F$ within $O((m+n)k^2)$ at the beginning. Thus this cost can be omitted in the complexity analysis. For other matrices in (14), \mathbf{P} and \mathbf{Q} are needed, but they have been pre-computed and cached in memory during computing $L^+(\theta)$. Thus all matrices in (14) can be obtained in $O((m+n)k^2)$ time and cached in $O(k^2)$ space. Then in (53), the Frobenius inner products between these matrices cost only $O(k^2)$ time. Through combining $L^+(\theta)$ calculation discussed in Section 3.2, Algorithm 2 summarizes the procedure of evaluating the objective function.

A.2 Detailed Derivation of a Stochastic Gradient Method Based on (46)

From (46), by defining $\bar{a}_i = \frac{a_i}{|\mathbb{O}_{\cdot, i}|}$ and $\bar{b}_j = \frac{b_j}{|\mathbb{O}_{\cdot, j}|}$, $L(\theta)$ in (10) can be written as

$$\lambda \Psi(\theta) + \sum_{(i,j) \in \mathbb{O}} \ell_{ij}^*(Y_{ij}, \hat{Y}_{ij}) + \omega \sum_{(i,j') \in \mathbb{O}} \sum_{(i',j) \in \mathbb{O}} \frac{1}{2} \bar{a}_i \bar{b}_j (\tilde{Y}_{ij} - \hat{Y}_{ij})^2. \tag{54}$$

Then $\nabla L^-(\theta)$ in (24) can be rewritten as

$$\nabla L^-(\theta) = \sum_{(i,j') \in \mathbb{O}} \sum_{(i',j) \in \mathbb{O}} \left[\frac{\partial f^i}{\partial \theta^u} \mathbf{q}_j \right] \bar{a}_i \bar{b}_j (\mathbf{p}_i^\top \mathbf{q}_j - \tilde{\mathbf{p}}_i^\top \tilde{\mathbf{q}}_j) \tag{55}$$

$$= \begin{bmatrix} \sum_{(i,j') \in \mathbb{O}} \frac{\partial f^i}{\partial \theta^u} \mathbf{Q}_0 \bar{a}_i \mathbf{p}_i - \hat{\mathbf{Q}}_0^\top \bar{a}_i \tilde{\mathbf{p}}_i \\ \sum_{(i',j) \in \mathbb{O}} \frac{\partial g^j}{\partial \theta^v} (\mathbf{P}_0 \bar{b}_j \mathbf{q}_j - \hat{\mathbf{P}}_0^\top \bar{b}_j \tilde{\mathbf{q}}_j) \end{bmatrix}, \tag{56}$$

where similar to (14), we define

$$\begin{aligned}
\mathbf{P}_0 &= \sum_{(i,j') \in \mathbb{O}} \bar{a}_i \mathbf{p}_i \mathbf{p}_i^\top, & \mathbf{Q}_0 &= \sum_{(i',j) \in \mathbb{O}} \bar{b}_j \mathbf{q}_j \mathbf{q}_j^\top, \\
\hat{\mathbf{P}}_0 &= \sum_{(i,j') \in \mathbb{O}} \bar{a}_i \tilde{\mathbf{p}}_i \tilde{\mathbf{p}}_i^\top, & \hat{\mathbf{Q}}_0 &= \sum_{(i',j) \in \mathbb{O}} \bar{b}_j \tilde{\mathbf{q}}_j \tilde{\mathbf{q}}_j^\top. \end{aligned} \tag{57}$$

and we omit details from (55) to (56) as the derivations are similar to those between (24) and (26).

On the other hand, from (17) and (18), we write $\nabla L^+(\theta)$ as

$$\nabla L^+(\theta) = \begin{bmatrix} \sum_{(i,j) \in \mathbb{O}} \frac{\partial f^i}{\partial \theta^u} \mathbf{q}_j X_{ij} \\ \sum_{(i,j) \in \mathbb{O}} \frac{\partial g^j}{\partial \theta^v} \mathbf{p}_i X_{ij} \end{bmatrix}. \tag{58}$$

Because both (56) and (58) are now operations summing over \mathbb{O} , $\nabla L(\theta)$ can be written as

$$\begin{bmatrix} \sum_{(i,j) \in \mathbb{O}} \frac{\partial f^i}{\partial \theta^u} (\mathbf{q}_j X_{ij} + \omega (\mathbf{Q}_0 \bar{a}_i \mathbf{p}_i - \hat{\mathbf{Q}}_0^\top \bar{a}_i \tilde{\mathbf{p}}_i)) \\ \sum_{(i,j) \in \mathbb{O}} \frac{\partial g^j}{\partial \theta^v} (\mathbf{p}_i X_{ij} + \omega (\mathbf{P}_0 \bar{b}_j \mathbf{q}_j - \hat{\mathbf{P}}_0^\top \bar{b}_j \tilde{\mathbf{q}}_j)) \end{bmatrix} + \lambda \nabla \Psi(\theta). \tag{59}$$

To derive a subsampled approach, we must consider two subsets $\hat{\mathbb{O}}_1 \subset \mathbb{O}$ and $\hat{\mathbb{O}}_2 \subset \mathbb{O}$ because from (57), two summations respectively over \mathbb{O} are involved in the above gradient. This results in the following estimate of $\nabla L(\theta)$.

$$\frac{|\mathbb{O}|}{|\hat{\mathbb{O}}_1|} \begin{bmatrix} \sum_{(i,j) \in \hat{\mathbb{O}}_1} \frac{\partial f^i}{\partial \theta^u} (\mathbf{q}_j X_{ij} + \omega (\mathbf{Q}_2 \bar{a}_i \mathbf{p}_i - \hat{\mathbf{Q}}_2^\top \bar{a}_i \tilde{\mathbf{p}}_i)) \\ \sum_{(i,j) \in \hat{\mathbb{O}}_1} \frac{\partial g^j}{\partial \theta^v} (\mathbf{p}_i X_{ij} + \omega (\mathbf{P}_2 \bar{b}_j \mathbf{q}_j - \hat{\mathbf{P}}_2^\top \bar{b}_j \tilde{\mathbf{q}}_j)) \end{bmatrix} + \lambda \nabla \Psi(\theta), \tag{60}$$

where

$$\begin{aligned}
\mathbf{P}_2 &= \frac{|\mathbb{O}|}{|\hat{\mathbb{O}}_2|} \sum_{(i,j) \in \hat{\mathbb{O}}_2} \bar{a}_i \mathbf{p}_i \mathbf{p}_i^\top, & \mathbf{Q}_2 &= \frac{|\mathbb{O}|}{|\hat{\mathbb{O}}_2|} \sum_{(i,j) \in \hat{\mathbb{O}}_2} \bar{b}_j \mathbf{q}_j \mathbf{q}_j^\top, \\
\hat{\mathbf{P}}_2 &= \frac{|\mathbb{O}|}{|\hat{\mathbb{O}}_2|} \sum_{(i,j) \in \hat{\mathbb{O}}_2} \bar{a}_i \tilde{\mathbf{p}}_i \tilde{\mathbf{p}}_i^\top, & \hat{\mathbf{Q}}_2 &= \frac{|\mathbb{O}|}{|\hat{\mathbb{O}}_2|} \sum_{(i,j) \in \hat{\mathbb{O}}_2} \bar{b}_j \tilde{\mathbf{q}}_j \tilde{\mathbf{q}}_j^\top. \end{aligned} \tag{61}$$

To make (60) an unbiased estimate of $\nabla L(\theta)$, we need that $\hat{\mathbb{O}}_1$ and $\hat{\mathbb{O}}_2$ are independent of each other, and have the scaling factors $\frac{|\mathbb{O}|}{|\hat{\mathbb{O}}_1|}$, $\frac{|\mathbb{O}|}{|\hat{\mathbb{O}}_2|}$ respectively in (60) and (61).

Because $\hat{\mathbb{O}}_1$ and $\hat{\mathbb{O}}_2$ are independent subsets, we can swap them in (60) to have another estimate of $\nabla L(\theta)$. Then four matrices $\mathbf{P}_1, \mathbf{Q}_1, \hat{\mathbf{P}}_1$ and $\hat{\mathbf{Q}}_1$ similar to those in (61) must be calculated, but the summations are now over $\hat{\mathbb{O}}_1$. In the end, the two subsampled gradients can be averaged.

Similar to the discussion on (14) and (27), for each pair of $\hat{\mathbb{O}}_1$ and $\hat{\mathbb{O}}_2$, (59) can be computed in

$$O((|\hat{\mathbb{O}}_1| + |\hat{\mathbb{O}}_2|)(k + k^2 + F(f) + F(g))) \tag{62}$$

time. We take $O(|\hat{\mathbb{O}}_2|k^2)$ term as an example to illustrate the difference from (27). Now for matrices in (61), each involves $O(|\hat{\mathbb{O}}_2|k^2)$ cost, but for those in (14), $O(mk^2)$ or $O(nk^2)$ are needed. Therefore, from (28) to (30), we can see the $m+n$ term should be replaced by $|\hat{\mathbb{O}}_1| + |\hat{\mathbb{O}}_2|$.

Next we check the complexity. We have mentioned in Section 4.1 that the task of going over all mn pairs is now replaced by sampling $\hat{\mathbb{O}}_1 \times \hat{\mathbb{O}}_2$ to cover $\mathbb{O} \times \mathbb{O}$. Thus the time complexity is by multiplying the $O(\frac{|\mathbb{O}|^2}{|\hat{\mathbb{O}}_1| |\hat{\mathbb{O}}_2|})$ steps and the cost in (62) for each SG step to have

$$O\left(\left(\frac{|\mathbb{O}|^2}{|\hat{\mathbb{O}}_1|} + \frac{|\mathbb{O}|^2}{|\hat{\mathbb{O}}_2|}\right)(k + k^2 + F(f) + F(g))\right). \tag{63}$$

Algorithm 4: An efficient implementation of Gauss-Newton method for solving (4)

Input: $\mathbf{U} = \{\dots, \mathbf{u}_i, \dots\}, \mathbf{V} = \{\dots, \mathbf{v}_j, \dots\}, \mathbb{O}, \tilde{\mathbf{P}}, \tilde{\mathbf{Q}}, \mathbf{A}, \mathbf{B}$.
Init:
 Draw $\boldsymbol{\theta}$ randomly
 $\langle \tilde{\mathbf{P}}_c, \tilde{\mathbf{Q}}_c \rangle_{\mathbb{F}} \leftarrow \langle \tilde{\mathbf{P}}^{\top} \mathbf{A} \tilde{\mathbf{P}}, \tilde{\mathbf{Q}}^{\top} \mathbf{B} \tilde{\mathbf{Q}} \rangle_{\mathbb{F}}$
 $L(\boldsymbol{\theta}), \mathbf{P}, \mathbf{Q}, \mathbf{P}_c, \mathbf{Q}_c, \hat{\mathbf{P}}_c, \hat{\mathbf{Q}}_c \leftarrow$
 Alg. 2($\boldsymbol{\theta}, \mathbf{U}, \mathbf{V}, \mathbb{O}, \tilde{\mathbf{P}}, \tilde{\mathbf{Q}}, \mathbf{A}, \mathbf{B}, \langle \tilde{\mathbf{P}}_c, \tilde{\mathbf{Q}}_c \rangle_{\mathbb{F}}$)

- 1 **repeat**
- 2 $\nabla L(\boldsymbol{\theta}) \leftarrow$ Alg. 1($\boldsymbol{\theta}, \mathbb{O}, \tilde{\mathbf{P}}, \tilde{\mathbf{Q}}, \mathbf{A}, \mathbf{B}, \mathbf{P}, \mathbf{Q}, \mathbf{P}_c, \mathbf{Q}_c, \hat{\mathbf{P}}_c, \hat{\mathbf{Q}}_c$)
- 3 $\mathbf{s} \leftarrow$ Alg. 3($\boldsymbol{\theta}, \mathbb{O}, \nabla L(\boldsymbol{\theta}), \mathbf{A}, \mathbf{B}, \mathbf{P}, \mathbf{Q}, \mathbf{P}_c, \mathbf{Q}_c$)
- 4 $\delta \leftarrow 1$
- 5 **while** *line-search steps are within a limit* **do**
- 6 $L(\boldsymbol{\theta} + \delta \mathbf{s}), \mathbf{P}, \mathbf{Q}, \mathbf{P}_c, \mathbf{Q}_c, \hat{\mathbf{P}}_c, \hat{\mathbf{Q}}_c \leftarrow$
 $\text{Alg. 2}(\boldsymbol{\theta} + \delta \mathbf{s}, \mathbf{U}, \mathbf{V}, \mathbb{O}, \tilde{\mathbf{P}}, \tilde{\mathbf{Q}}, \mathbf{A}, \mathbf{B}, \langle \tilde{\mathbf{P}}_c, \tilde{\mathbf{Q}}_c \rangle_{\mathbb{F}})$
- 7 **if** $L(\boldsymbol{\theta} + \delta \mathbf{s}) \leq L(\boldsymbol{\theta}) + \eta \delta \mathbf{s}^{\top} \nabla L(\boldsymbol{\theta})$ **then break**
- 8 $\delta \leftarrow \delta/2$
- 9 **end**
- 10 $\boldsymbol{\theta} \leftarrow \boldsymbol{\theta} + \delta \mathbf{s}$
- 11 **until** *stopping condition is satisfied*

Algorithm 2: Objective function evaluation: $L(\boldsymbol{\theta})$

Input: $\boldsymbol{\theta}, \mathbf{U}, \mathbf{V}, \mathbb{O}, \tilde{\mathbf{P}}, \tilde{\mathbf{Q}}, \mathbf{A}, \mathbf{B}, \langle \tilde{\mathbf{P}}_c, \tilde{\mathbf{Q}}_c \rangle_{\mathbb{F}}$.
 1 $\mathbf{P} \leftarrow [\dots f(\boldsymbol{\theta}^{\mathbf{u}}; \mathbf{u}_i) \dots]^{\top}, \mathbf{Q} \leftarrow [\dots g(\boldsymbol{\theta}^{\mathbf{v}}; \mathbf{v}_j) \dots]^{\top}$
 2 $\mathbf{P}_c \leftarrow \mathbf{P}^{\top} \mathbf{A} \mathbf{P}, \mathbf{Q}_c \leftarrow \mathbf{Q}^{\top} \mathbf{B} \mathbf{Q}$
 3 $\hat{\mathbf{P}}_c \leftarrow \tilde{\mathbf{P}}^{\top} \mathbf{A} \tilde{\mathbf{P}}, \hat{\mathbf{Q}}_c \leftarrow \tilde{\mathbf{Q}}^{\top} \mathbf{B} \tilde{\mathbf{Q}}$
 4 $L^+(\boldsymbol{\theta}) \leftarrow \sum_{(i,j) \in \mathbb{O}} \ell_{ij}^+(Y_{ij}, \mathbf{p}_i^{\top} \mathbf{q}_j)$
 5 $L^-(\boldsymbol{\theta}) \leftarrow \frac{1}{2} \langle \tilde{\mathbf{P}}_c, \tilde{\mathbf{Q}}_c \rangle_{\mathbb{F}} - \langle \hat{\mathbf{P}}_c, \hat{\mathbf{Q}}_c \rangle_{\mathbb{F}} + \frac{1}{2} \langle \mathbf{P}_c, \mathbf{Q}_c \rangle_{\mathbb{F}}$
 6 $L(\boldsymbol{\theta}) \leftarrow L^+(\boldsymbol{\theta}) + \omega L^-(\boldsymbol{\theta}) + \lambda \Psi(\boldsymbol{\theta})$
Output: $L(\boldsymbol{\theta}), \mathbf{P}, \mathbf{Q}, \mathbf{P}_c, \mathbf{Q}_c, \hat{\mathbf{P}}_c, \hat{\mathbf{Q}}_c$

Algorithm 3: Conjugate gradient (CG) procedures

Input: $\boldsymbol{\theta}, \mathbb{O}, \nabla L(\boldsymbol{\theta}), \mathbf{A}, \mathbf{B}, \mathbf{P}, \mathbf{Q}, \mathbf{P}_c, \mathbf{Q}_c$.
 1 $\mathbf{r} \leftarrow -\nabla L(\boldsymbol{\theta}), \mathbf{d} \leftarrow \mathbf{r}, \mathbf{s} \leftarrow \mathbf{0}, \gamma \leftarrow \mathbf{r}^{\top} \mathbf{r}, \gamma_0 \leftarrow \gamma$
 2 **repeat**
 3 $\mathbf{W} \leftarrow [\dots \frac{\partial f^i}{\partial \theta^{\mathbf{u}}} \mathbf{d}^{\mathbf{u}} \dots]^{\top}, \mathbf{H} \leftarrow [\dots \frac{\partial g^j}{\partial \theta^{\mathbf{v}}} \mathbf{d}^{\mathbf{v}} \dots]^{\top}$
 4 $\mathbf{W}_c \leftarrow \mathbf{W}^{\top} \mathbf{A} \mathbf{P}, \mathbf{H}_c \leftarrow \mathbf{H}^{\top} \mathbf{B} \mathbf{Q}$
 5 $Z_{ij} \leftarrow \frac{\partial^2 \ell_{ij}^+}{\partial Y_{ij}^2} (\mathbf{w}_i^{\top} \mathbf{q}_j + \mathbf{p}_i^{\top} \mathbf{h}_j), \forall (i, j) \in \mathbb{O}$
 6 $\mathbf{Z}_q \leftarrow \mathbf{Z} \mathbf{Q}, \mathbf{Z}_p \leftarrow \mathbf{Z}^{\top} \mathbf{P}$
 7 $\mathbf{G} \mathbf{d} \leftarrow \left[\begin{array}{l} \langle \mathbf{J}^{\mathbf{u}}, \mathbf{Z}_q + \omega(\mathbf{A} \mathbf{W} \mathbf{Q}_c + \mathbf{A} \mathbf{P} \mathbf{H}_c) \rangle \\ \langle \mathbf{J}^{\mathbf{v}}, \mathbf{Z}_p + \omega(\mathbf{B} \mathbf{H} \mathbf{P}_c + \mathbf{B} \mathbf{Q} \mathbf{W}_c) \rangle \end{array} \right] + \lambda \nabla^2 \Psi(\boldsymbol{\theta}) \mathbf{d}$
 8 $\mathbf{t} \leftarrow \gamma / \mathbf{d}^{\top} \mathbf{G} \mathbf{d}$
 9 $\mathbf{s} \leftarrow \mathbf{s} + \mathbf{t} \mathbf{d}$
 10 $\mathbf{r} \leftarrow \mathbf{r} - \mathbf{t} \mathbf{G} \mathbf{d}$
 11 $\gamma_{\text{new}} = \mathbf{r}^{\top} \mathbf{r}$
 12 $\beta \leftarrow \gamma_{\text{new}} / \gamma$
 13 $\mathbf{d} \leftarrow \mathbf{r} + \beta \mathbf{d}$
 14 $\gamma \leftarrow \gamma_{\text{new}}$
 15 **until** $\sqrt{\gamma} \leq \xi \sqrt{\gamma_0}$ or the max number of CG steps is reached
Output: \mathbf{s}

A.3 Detailed Derivation of the SOGram Method

We discuss the Stochastic Online Gramian (SOGram) method proposed in [13], which is very related to the method discussed in Section A.2. It considers a special case of the objective function in (54) by excluding \bar{a}_i and \bar{b}_j . From this we show that SOGram is an extension.

In [13], they view matrices in (61) as estimates of matrices in (14). The main idea behind SOGram is to apply a variance reduction scheme to improve these estimates. Specifically, they maintain four zero-initialized matrices with the following exponentially averaged updates before each step

$$\begin{aligned} \mathbf{P}'_c &\leftarrow (1 - \alpha) \mathbf{P}'_c + \alpha \mathbf{P}_2, & \mathbf{Q}'_c &\leftarrow (1 - \alpha) \mathbf{Q}'_c + \alpha \mathbf{Q}_2, \\ \hat{\mathbf{P}}'_c &\leftarrow (1 - \alpha) \hat{\mathbf{P}}'_c + \alpha \hat{\mathbf{P}}_2, & \hat{\mathbf{Q}}'_c &\leftarrow (1 - \alpha) \hat{\mathbf{Q}}'_c + \alpha \hat{\mathbf{Q}}_2, \end{aligned} \quad (64)$$

where α is a hyper-parameter for controlling the bias-variance tradeoff. Then they replace (60) with

$$\frac{|\mathbb{O}|}{|\hat{\mathbb{O}}_1|} \left[\begin{array}{l} \sum_{(i,j) \in \hat{\mathbb{O}}_1} \frac{\partial f^i}{\partial \theta^{\mathbf{u}}} \left(\mathbf{q}_j X_{ij} + \omega(\mathbf{Q}'_c \bar{a}_i \mathbf{p}_i - \hat{\mathbf{Q}}_c^{\top} \bar{a}_i \hat{\mathbf{p}}_i) \right) \\ \sum_{(i,j) \in \hat{\mathbb{O}}_1} \frac{\partial g^j}{\partial \theta^{\mathbf{v}}} \left(\mathbf{p}_i X_{ij} + \omega(\mathbf{P}'_c \bar{b}_j \mathbf{q}_j - \hat{\mathbf{P}}_c^{\top} \bar{b}_j \hat{\mathbf{q}}_j) \right) \end{array} \right] + \lambda \nabla \Psi(\boldsymbol{\theta}), \quad (65)$$

which is considered as the subsampled gradient vector for each SG step.

A.4 Detailed Procedures of Gauss-Newton Method

Details of the CG method involving a sequence of $\mathbf{G} \mathbf{d}$ operations are given in Algorithm 3. The overall procedures of the Gauss-Newton method are summarized in Algorithm 4.