

PRACTICAL ALGORITHMS FOR LEARNING NEAR-ISOMETRIC LINEAR EMBEDDINGS*

JERRY LUO[†], KAYLA SHAPIRO[‡], HAO-JUN MICHAEL SHI[§], QI YANG[¶], AND KAN ZHU^{||}

Abstract. We propose two practical non-convex approaches for learning near-isometric, linear embeddings of finite sets of data points. Given a set of training points \mathcal{X} , we consider the secant set $S(\mathcal{X})$ that consists of all pairwise difference vectors of \mathcal{X} , normalized to lie on the unit sphere. The problem can be formulated as finding a symmetric and positive semi-definite matrix Ψ that preserves the norms of all the vectors in $S(\mathcal{X})$ up to a distortion parameter δ . Motivated by non-negative matrix factorization, we reformulate our problem into a Frobenius norm minimization problem, which is solved by the Alternating Direction Method of Multipliers (ADMM) and develop an algorithm, *FroMax*. Another method solves for a projection matrix Ψ by minimizing the restricted isometry property (RIP) directly over the set of symmetric, positive semi-definite matrices. Applying ADMM and a Moreau decomposition on a proximal mapping, we develop another algorithm, *NILE-Pro*, for dimensionality reduction. *FroMax* is shown to converge faster for smaller δ while *NILE-Pro* converges faster for larger δ . Both non-convex approaches are then empirically demonstrated to be more computationally efficient than prior convex approaches for a number of applications in machine learning and signal processing.

Key words. dimensionality reduction, linear embeddings, compressive sensing, approximate nearest neighbors, classification

1. Introduction.

1.1. Motivation. The size of the acquired raw data currently poses a challenge to current state-of-the-art information processing systems. Since many machine learning algorithms' computational efficiency scale with the complexity of the data, machine learning researchers have introduced a family of algorithms for *dimensionality reduction* to address this issue. Dimensionality reduction algorithms construct a concise representation of high-dimensional data on a lower-dimensional subspace, with as minimal loss of intrinsic information as possible. This representation is often referred to as a low-dimensional *embedding*.

The canonical approach in statistics for constructing a linear embedding is principal components analysis (PCA) [14]. PCA is a linear embedding technique that projects data points onto a lower-dimensional subspace spanned by the principal components that contain the most variability within the data. PCA benefits from being computationally efficient and easily generalizable to new data sets; however, it fails to preserve pairwise distances between sample data points, sometimes rendering two distinct points indistinguishable in the low-dimensional embedding space. This can potentially hamper the performance of PCA and other similar algorithms.

*This research was conducted as part of the California Research Training Program in Computational and Applied Mathematics 2014.

[†]Department of Mathematics, University of Arizona, Tucson, AZ 85721. (jerry-luo@math.arizona.edu)

[‡]Department of Computing, Imperial College London, London SW7 2AZ, United Kingdom. (k.shapiro@berkeley.edu)

[§]Department of Mathematics, University of California, Los Angeles, Los Angeles, CA 90024. (hjmshi@ucla.edu)

[¶]Department of Mathematics, University of Southern California, Los Angeles, 90089. (yangq@usc.edu)

^{||}Department of Computer Science, Columbia University, New York, NY 10027. (kzhu9@ucla.edu)

Other popular nonlinear, manifold learning methods, such as ISOMAP and locally linear embedding (LLE), preserve geometric structure by approximating geodesics from k -nearest neighbors. Although these methods produce embeddings which are easy to explicitly store and generalize, most fail to preserve all pairwise distances between data points.

A linear embedding technique that preserves all pairwise distances is the method of *random projections*. Given \mathcal{X} , a cloud of Q data points in a high-dimensional Euclidean space \mathbb{R}^N , the Johnson-Lindenstrauss Lemma [8] states that there exists a linear, near-isometric, or distance preserving, embedding such that \mathcal{X} can be mapped to a subspace of dimension $M = \mathcal{O}(\log Q)$ with high probability. Despite its conceptual simplicity, random projections suffers from probabilistic and asymptotic theoretical guarantees in Johnson-Lindenstrauss [8]. A random projections mapping is also independent of the data under consideration, failing to utilize the geometric structure of the data.

We would like to deterministically devise a linear embedding for dimensionality reduction that preserves all normalized pairwise distances. In particular, linear embeddings can be explicitly stored using a matrix operator and can therefore be quickly applied to any new data point.

1.2. Related Work. Using the geometric structure of the data, Hegde, et. al. developed a new deterministic approach, NuMax, to construct a near-isometric, linear embedding [6]. Given a training set $\mathcal{X} \subset \mathbb{R}^N$, the *secant set* is constructed by taking all pairwise difference vectors of \mathcal{X} , which are then normalized to lie on the unit sphere. Hegde, et. al. formulated a rank minimization problem with affine constraints to construct a *projection matrix* ψ that preserves norms of all vectors in $S(\mathcal{X})$ up to a distortion parameter δ . They then relax this problem to a convex program that can be solved using a tractable semidefinite program (SDP), with column generation for large data sets, and develop NuMax based on the Alternating Direction Method of Multipliers (ADMM). This framework deterministically produces a near-isometric linear embedding. Other algorithmic approaches for finding near-isometric linear embeddings are also described in [4, 5, 19].

1.3. Organization. The rest of the paper is organized as follows. We review the restricted isometry property and NuMax algorithm in §2. The FroMax algorithm is introduced in §3. NILE-Pro is discussed in §4. Rank adjustment and column generation methods which increase computational efficiency for large data sets is introduced in §5. Numerical simulations and runtime performance results are presented in §7. Lastly, §8 concludes the paper and gives direction for future work.

2. Background.

2.1. Restricted Isometry Property (RIP). E. Candes, et. al. introduced a formal, relaxed notion of isometry in [1] as follows:

DEFINITION 2.1. *Suppose $M \leq N$ and consider $\mathcal{X} \subset \mathbb{R}^N$. An embedding operator $\mathcal{P} : \mathcal{X} \rightarrow \mathbb{R}^M$ satisfies the restricted isometry property (RIP) on \mathcal{X} if there exists a positive constant $\delta > 0$ such that, for every $\mathbf{x}, \mathbf{x}' \in \mathcal{X}$, the following relation holds:*

$$(2.1) \quad (1 - \delta)\|\mathbf{x} - \mathbf{x}'\|_2^2 \leq \|\mathcal{P}\mathbf{x} - \mathcal{P}\mathbf{x}'\|_2^2 \leq (1 + \delta)\|\mathbf{x} - \mathbf{x}'\|_2^2.$$

We may also refer to δ as the *isometry constant*. Intuitively, this notion of near-isometry requires the distance of every pair of points in \mathcal{X} to be almost preserved.

Hegde, et. al. [6] leveraged this idea to develop a framework for finding low rank matrices that satisfy the RIP.

2.2. NuMax. In this section, we review Hegde et. al. [6]’s work on NuMax. Given a data set $\mathcal{X} \subset \mathbb{R}^N$, Hegde et. al. formulate the secant set as follows:

$$(2.2) \quad S(\mathcal{X}) = \left\{ \frac{\mathbf{x} - \mathbf{x}'}{\|\mathbf{x} - \mathbf{x}'\|_2}, \mathbf{x}, \mathbf{x}' \in \mathcal{X}, \mathbf{x} \neq \mathbf{x}' \right\}.$$

Hegde, et. al. seeks to find a projection matrix $\Psi \in \mathbb{R}^{M \times N}$ with the smallest possible rank that satisfies the RIP on $S(\mathcal{X})$ for a given $\delta > 0$. This problem is then cast as an optimization problem over all symmetric matrices which we denote as $\mathbb{S}^{N \times N}$. Let $\mathbf{P} = \Psi^T \Psi \in \mathbb{S}^{N \times N}$ with $\text{rank}(\mathbf{P}) = M$. Then for all secants $\mathbf{v}_i \in S(\mathcal{X})$, we may rewrite the RIP constraint as:

$$(2.3) \quad (1 - \delta)\|\mathbf{v}_i\|_2^2 \leq \|\Psi \mathbf{v}_i\|_2^2 \leq (1 + \delta)\|\mathbf{v}_i\|_2^2.$$

After grouping we get:

$$(2.4) \quad \left| \|\Psi \mathbf{v}_i\|_2^2 - \|\mathbf{v}_i\|_2^2 \right| \leq \delta \|\mathbf{v}_i\|_2^2.$$

Then we isolate δ and use the fact that $\|\Psi \mathbf{v}_i\|_2^2 = \mathbf{v}_i^T \mathbf{P} \mathbf{v}_i$:

$$(2.5) \quad \left| \|\Psi \mathbf{v}_i\|_2^2 - 1 \right| \leq \delta$$

$$(2.6) \quad \left| \mathbf{v}_i^T \mathbf{P} \mathbf{v}_i - 1 \right| \leq \delta.$$

Let $\mathbf{1}_S$ denote the S -dimensional ones vector and $\mathcal{A} : \mathbb{R}^{N \times N} \rightarrow \mathbb{R}^S$ where $S = |S(\mathcal{X})|$, and $\mathcal{A} : X \mapsto \{\mathbf{v}_i^T X \mathbf{v}_i\}_{i=1}^S$, where the output is understood to be an S dimensional vector with the i -th entry being $\mathbf{v}_i^T X \mathbf{v}_i$. This admits the rank minimization problem:

$$(2.7) \quad \begin{aligned} & \underset{\mathbf{P}}{\text{minimize}} && \text{rank}(\mathbf{P}) \\ & \text{subject to} && \|\mathcal{A}(\mathbf{P}) - \mathbf{1}_S\|_\infty \leq \delta \\ & && \mathbf{P} \succeq 0. \end{aligned}$$

Here, $\mathbf{P} \succeq 0$ means \mathbf{P} is symmetric positive semidefinite. However, since rank minimization is a non-convex, NP-hard problem [16], a convex relaxation is performed on the objective to obtain the following nuclear-norm minimization:

$$(2.8) \quad \begin{aligned} & \underset{\mathbf{P}}{\text{minimize}} && \|\mathbf{P}\|_* \\ & \text{subject to} && \|\mathcal{A}(\mathbf{P}) - \mathbf{1}_S\|_\infty \leq \delta \\ & && \mathbf{P} \succeq 0, \end{aligned}$$

where $\|\mathbf{P}\|_*$ is the nuclear norm of \mathbf{P} , or the sum of the singular values of \mathbf{P} . Then the desired linear embedding $\Psi \in \mathbb{R}^{M \times N}$ can be found by taking a matrix square root of the minimizer $\mathbf{P}^* = \mathbf{U} \mathbf{\Gamma} \mathbf{U}^T$ by

$$(2.9) \quad \Psi = \mathbf{\Gamma}_M^{1/2} \mathbf{U}_M^T,$$

where $\mathbf{\Gamma}_M = \text{diag}\{\lambda_1, \dots, \lambda_M\}$ denotes the M leading (non-zero) eigenvalues of \mathbf{P}^* and \mathbf{U}_M are the corresponding eigenvectors.

Applying the Alternating Direction Method of Multipliers (ADMM), the optimization problem is rewritten by introducing auxiliary variables $\mathbf{L} \in \mathbb{S}^{N \times N}$ and $\mathbf{q} \in \mathbb{R}^S$:

$$\begin{aligned}
(2.10) \quad & \underset{\mathbf{P}, \mathbf{L}, \mathbf{q}}{\text{minimize}} && \|\mathbf{P}\|_* \\
& \text{subject to} && \mathbf{P} = \mathbf{L} \\
& && \mathcal{A}(\mathbf{L}) = \mathbf{q}, \\
& && \|\mathbf{q} - \mathbf{1}_S\|_\infty \leq \delta \\
& && \mathbf{P} \succeq 0.
\end{aligned}$$

Introducing auxiliary variables \mathbf{L} and \mathbf{q} allows constrained optimization techniques to be applied while breaking up the problem into simpler subproblems. The linear constraints are then relaxed to form an augmented Lagrangian as follows:

$$(2.11) \quad L_A(\mathbf{P}, \mathbf{L}, \mathbf{q}; \mathbf{\Gamma}, \boldsymbol{\omega}) = \|\mathbf{P}\|_* + \frac{\beta_1}{2} \|\mathbf{P} - \mathbf{L} - \mathbf{\Gamma}\|_F^2 + \frac{\beta_2}{2} \|\mathcal{A}(\mathbf{L}) - \mathbf{q} - \boldsymbol{\omega}\|_2^2,$$

where $\mathbf{\Gamma} \in \mathbb{S}^{N \times N}$ and $\boldsymbol{\omega} \in \mathbb{R}^S$ represent the scaled Lagrange multipliers and $\beta_1, \beta_2 \in \mathbb{R}$ are chosen parameters.

NuMax then solves the following augmented Lagrangian problem:

$$\begin{aligned}
(2.12) \quad & \underset{\mathbf{P}, \mathbf{L}, \mathbf{q}, \mathbf{\Gamma}, \boldsymbol{\omega}}{\text{minimize}} && L_A(\mathbf{P}, \mathbf{L}, \mathbf{q}, \mathbf{\Gamma}, \boldsymbol{\omega}) \\
& \text{subject to} && \|\mathbf{q} - \mathbf{1}_S\|_\infty \leq \delta \\
& && \mathbf{P} \succeq 0,
\end{aligned}$$

\mathbf{P}, \mathbf{L} and \mathbf{q} are optimized in an alternating fashion, i.e. optimized one at a time with the others held fixed. This optimization can then be solved by three easier sub-problems, admitting a computationally efficient solution.

For more information regarding theoretical and empirical properties of NuMax, please refer to Hegde et. al. [6].

This framework, though slower than conventional methods such as PCA and random projections, deterministically admits a projection matrix satisfying the RIP with low rank. However, NuMax computes a singular value decomposition of \mathbf{P} each iteration, which is computationally expensive. Furthermore, though minimizing the nuclear-norm tends to give low rank matrices, NuMax does not theoretically guarantee the lowest rank embedding for a given δ .

These issues motivate the pursuit of other practical algorithms that optimize similar non-convex problems that may admit low rank, near-isometric projection matrices that give faster, but sufficient (not necessarily optimal) results. Rather than solving both the rank minimization and near-isometry problems simultaneously, we solve a simpler non-convex problem quickly to find a near-isometric projection matrix and apply a rank adjustment heuristic to choose a minimal rank.

3. FroMax. Our first algorithm, *Frobenius norm minimization with Max-norm constraints*, or *FroMax*, combines ideas from NuMax and NMF to formulate a Frobenius norm minimization problem which we then solve based on ADMM, similar to NuMax [22]. Though this algorithm does not discover the optimal rank for $\boldsymbol{\Psi}$, we combine FroMax with a rank adjustment heuristic to find low rank embeddings.

3.1. Non-Negative Matrix Factorization (NMF). FroMax is motivated by ideas from *non-negative matrix factorization*. Non-negative matrix factorization (NMF) is a group of algorithms that factorize a non-negative matrix \mathbf{V} into two low-rank non-negative matrices \mathbf{W} and \mathbf{H} [12]. More rigorously, let $\mathbf{V} \in \mathbb{R}^{N \times M}$ be given, then we solve for $\mathbf{W} \in \mathbb{R}^{M \times Q}$, and $\mathbf{H} \in \mathbb{R}^{Q \times N}$ by solving the following optimization problem:

$$(3.1) \quad \begin{aligned} & \underset{\mathbf{W}, \mathbf{H}}{\text{minimize}} && \|\mathbf{WH} - \mathbf{V}\|_F^2 \\ & \text{subject to} && W_{ij} \geq 0, H_{ij} \geq 0, \forall i, j. \end{aligned}$$

NMF motivates the problem formulation for our first algorithm, FroMax.

3.2. Optimization Framework. We formulate a specialized matrix factorization minimization problem to solve for a near-isometric linear embedding as follows: Given a desired rank $r \ll N$, let $\Psi \in \mathbb{R}^{r \times N}$ be the variable projection matrix. Here, we seek to solve:

$$(3.2) \quad \begin{aligned} & \underset{\mathbf{P}, \Psi}{\text{minimize}} && \frac{1}{2} \|\mathbf{P} - \Psi^T \Psi\|_F^2 \\ & \text{subject to} && \|\mathcal{A}(\mathbf{P}) - \mathbf{1}_S\|_\infty \leq \delta, \end{aligned}$$

where $\mathbf{P} \in \mathbb{R}^{N \times N}$ is a variable.

We introduce auxiliary variables to simplify our optimization and break up our problem into simpler subproblems that can be solved by applying ADMM. In particular, let $\mathbf{Y} = \Psi \in \mathbb{R}^{r \times N}$, $\mathbf{X} = \mathbf{Y}^T \in \mathbb{R}^{N \times r}$. Then

$$(3.3) \quad \begin{aligned} & \underset{\mathbf{P}, \mathbf{X}, \mathbf{Y}, \mathbf{q}}{\text{minimize}} && \frac{1}{2} \|\mathbf{P} - \mathbf{X}\mathbf{Y}\|_F^2 \\ & \text{subject to} && \mathcal{A}(\mathbf{P}) = \mathbf{q} \\ & && \mathbf{Y} = \mathbf{X}^T \\ & && \|\mathbf{q} - \mathbf{1}_S\|_\infty \leq \delta. \end{aligned}$$

This gives $\mathbf{Y} = \Psi \in \mathbb{R}^{r \times N}$ such that the RIP holds for all secant vectors in the secant set $S(\mathcal{X})$ for an isometry constant δ . The optimization formulation for (3.2) is conceptually simple, only requiring the input data set \mathcal{X} , desired isometry constant $\delta > 0$ and desired rank r .

An important caveat is that our optimization problem is non-convex. Thus, we cannot guarantee that FroMax will converge to the optimal solution of (3.2). However, various experiments in §7 indicate that FroMax yields excellent, stable results for real-world data sets and finds projection matrices faster than NuMax and other convex approaches. We implement ADMM since Wang et. al. [21] indicate that ADMM is more likely to converge than the Augmented Lagrangian Method for nonconvex, nonsmooth problems.

3.3. ADMM. We develop our algorithm, FroMax, to solve (3.3) based on ADMM. We relax the linear constraints and form an augmented Lagrangian of (3.3) as follows:

$$(3.4) \quad \begin{aligned} L_A(\mathbf{X}, \mathbf{Y}, \mathbf{q}, \mathbf{P}) = & \frac{1}{2} \|\mathbf{P} - \mathbf{X}\mathbf{Y}\|_F^2 + \lambda^T (\mathcal{A}(\mathbf{P}) - \mathbf{q}) + \langle \Gamma, \mathbf{Y} - \mathbf{X}^T \rangle \\ & + \frac{\beta_1}{2} \|\mathcal{A}(\mathbf{P}) - \mathbf{q}\|_2^2 + \frac{\beta_2}{2} \|\mathbf{Y} - \mathbf{X}^T\|_F^2 + \iota_{\{\mathbf{q}; \|\mathbf{q} - \mathbf{1}_S\|_\infty \leq \delta\}}. \end{aligned}$$

Here, $\boldsymbol{\lambda} \in \mathbb{R}^N$ and $\boldsymbol{\Gamma} \in \mathbb{R}^{r \times N}$ represent the scaled Lagrange multipliers. The indicator function, $\iota_{\{\mathbf{q}: \|\mathbf{q} - \mathbf{1}_S\|_\infty \leq \delta\}}$, is defined as

$$\iota_{\{\mathbf{q}: \|\mathbf{q} - \mathbf{1}_S\|_\infty \leq \delta\}} = \begin{cases} 0 & \text{if } \|\mathbf{q} - \mathbf{1}_S\|_\infty \leq \delta \\ \infty & \text{otherwise.} \end{cases}$$

The optimization in (3.4) is carried out over $\mathbf{P} \in \mathbb{R}^{N \times N}$, $\mathbf{X} \in \mathbb{R}^{N \times r}$, $\mathbf{Y} \in \mathbb{R}^{r \times N}$, and $\mathbf{q} \in \mathbb{R}^S$, while $\boldsymbol{\lambda}$ and $\boldsymbol{\Gamma}$ are also iteratively updated. We optimize each variable in an alternating fashion like NuMax. The following steps below are performed until convergence.

Update \mathbf{q} : Isolating the terms that involve \mathbf{q} , we obtain a new estimate \mathbf{q}_{k+1} as the solution of the constrained optimization problem

$$(3.5) \quad \mathbf{q}_{k+1} \leftarrow \arg \min_{\mathbf{q}} \boldsymbol{\lambda}_k^T (\mathcal{A}(\mathbf{P}_k) - \mathbf{q}) + \frac{\beta_1}{2} \|\mathcal{A}(\mathbf{P}_k) - \mathbf{q}\|_2^2 + \iota_{\{\mathbf{q}: \|\mathbf{q} - \mathbf{1}_S\|_\infty \leq \delta\}}.$$

Define $\mathbf{z} = \mathcal{A}(\mathbf{P}_k) - \boldsymbol{\lambda}_k - \mathbf{1}_S$. Using a component-wise truncation procedure for entries in \mathbf{q} , we easily see that

$$(3.6) \quad \mathbf{q}_{k+1} = \mathbf{1}_S + \text{sign}(\mathbf{z}) \cdot \min(|\mathbf{z}|, \delta),$$

where the sign and min operators are applied component-wise.

Update \mathbf{P} : Isolating the terms that involve \mathbf{P} , we obtain a new estimate \mathbf{P}_{k+1} as the solution of the constrained optimization problem

$$(3.7) \quad \mathbf{P}_{k+1} \leftarrow \arg \min_{\mathbf{P}} \frac{1}{2} \|\mathbf{P} - \mathbf{X}_k \mathbf{Y}_k\|_F^2 + \boldsymbol{\lambda}_k^T (\mathcal{A}(\mathbf{P}) - \mathbf{q}_{k+1}) + \frac{\beta_1}{2} \|\mathcal{A}(\mathbf{P}) - \mathbf{q}_{k+1}\|_2^2,$$

such that $\mathbf{P} \succeq 0$. Since this is a least-squares problem, we can solve for the minimum by solving:

$$(3.8) \quad (\mathbf{P} - \mathbf{X}_k \mathbf{Y}_k) + \sum_{j=1}^S \lambda_{k,j} \mathbf{v}_j \mathbf{v}_j^T + \beta_1 \mathcal{A}^* (\mathcal{A}(\mathbf{P}) - \mathbf{q}_{k+1}) = 0,$$

where \mathcal{A}^* is the adjoint of \mathcal{A} . The solution is

$$(3.9) \quad \mathbf{P}_{k+1} = (\mathbf{I} + \beta_1 \mathcal{A}^* \mathcal{A})^\dagger (\beta_1 \mathcal{A}^* \mathbf{q}_{k+1} + \mathbf{X}_k \mathbf{Y}_k - \sum_{j=1}^S \lambda_{k,j} \mathbf{v}_j \mathbf{v}_j^T)$$

where \dagger denotes the pseudoinverse.

Update \mathbf{X} : Isolating the terms that involve \mathbf{X} , we obtain a new estimate \mathbf{X}_{k+1} as the solution of the constrained optimization problem

$$(3.10) \quad \mathbf{X}_{k+1} \leftarrow \arg \min_{\mathbf{X}} \frac{1}{2} \|\mathbf{P}_{k+1} - \mathbf{X} \mathbf{Y}_k\|_F^2 + \langle \boldsymbol{\Gamma}_k, \mathbf{Y}_k - \mathbf{X}^T \rangle + \frac{\beta_2}{2} \|\mathbf{Y}_k - \mathbf{X}^T\|_F^2.$$

It is easily seen that this can be solved similarly to the \mathbf{P} update.

Update \mathbf{Y} : Isolating the terms that involve \mathbf{Y} , we obtain a new estimate \mathbf{Y}_{k+1} as the solution of the constrained optimization problem

$$(3.11) \quad \mathbf{Y}_{k+1} \leftarrow \arg \min_{\mathbf{Y}} \frac{1}{2} \|\mathbf{P}_{k+1} - \mathbf{X}_{k+1} \mathbf{Y}\|_F^2 + \langle \boldsymbol{\Gamma}_k, \mathbf{Y} - \mathbf{X}_{k+1}^T \rangle + \frac{\beta_2}{2} \|\mathbf{Y} - \mathbf{X}_{k+1}^T\|_F^2.$$

Algorithm 1 FroMax

Inputs: Secant set $\mathcal{S}(\mathcal{X}) = \{\mathbf{v}_i\}_{i=1}^S$, isometry constant δ , desired rank for \mathbf{P} r , max iterations $m > 0$

Parameters: $\beta_1, \beta_2, \eta > 0$

Output: $\mathbf{P}, \mathbf{X}, \mathbf{Y}$

for $k = 1$ **to** m **do**

$\mathbf{z} \leftarrow \mathcal{A}(\mathbf{P}_k) - \boldsymbol{\lambda}_k - \mathbf{1}_S$

$\mathbf{q}_{k+1} \leftarrow \mathbf{1}_S + \text{sign}(\mathbf{z}) \cdot \min(|\mathbf{z}|, \delta)$

$\mathbf{P}_{k+1} \leftarrow (\mathbf{I} + \beta_1 \mathcal{A}^* \mathcal{A})^\dagger (\beta_1 \mathcal{A}^* \mathbf{q}_{k+1} + \mathbf{X}_k \mathbf{Y}_k - \sum_{j=1}^S \lambda_{k,j} \mathbf{v}_j \mathbf{v}_j^T)$

$\mathbf{X}_{k+1} \leftarrow (\boldsymbol{\Gamma}_k^T + \beta_2 \mathbf{Y}_k^T + \mathbf{P}_{k+1} \mathbf{Y}_k^T) (\mathbf{Y}_k \mathbf{Y}_k^T + \beta_2 \mathbf{I})^{-1}$

$\mathbf{Y}_{k+1} \leftarrow (\mathbf{X}_{k+1}^T \mathbf{X}_{k+1} + \beta_2 \mathbf{I})^{-1} (\mathbf{X}_{k+1}^T \mathbf{P}_{k+1} - \boldsymbol{\Gamma}_k + \beta_2 \mathbf{X}_{k+1}^T)$

$\boldsymbol{\lambda}_{k+1} \leftarrow \boldsymbol{\lambda}_k + \eta \beta_1 (\mathcal{A}(\mathbf{P}_{k+1}) - \mathbf{q}_{k+1})$

$\boldsymbol{\Gamma}_{k+1} \leftarrow \boldsymbol{\Gamma}_k + \eta \beta_2 (\mathbf{Y}_{k+1} - \mathbf{X}_{k+1}^T)$

if $\frac{1}{2} \|\mathbf{P}_{k+1} - \mathbf{X}_{k+1} \mathbf{Y}_{k+1}\|_F^2 < \epsilon$ **then**
 break

end if

end for

return $\mathbf{P} \leftarrow \mathbf{P}_k, \mathbf{X} \leftarrow \mathbf{X}_k, \mathbf{Y} \leftarrow \mathbf{Y}_k$

It is easily seen that this can be solved similarly to the \mathbf{X} update.

Update $\boldsymbol{\lambda}, \boldsymbol{\Gamma}$: Following standard augmented Lagrangian methods, we update $\boldsymbol{\Gamma}, \boldsymbol{\Pi}$ according to the following equations

$$(3.12) \quad \boldsymbol{\lambda}_{k+1} \leftarrow \boldsymbol{\lambda}_k + \eta \beta_1 (\mathcal{A}(\mathbf{P}_{k+1}) - \mathbf{q}_{k+1})$$

$$(3.13) \quad \boldsymbol{\Gamma}_{k+1} \leftarrow \boldsymbol{\Gamma}_k + \eta \beta_2 (\mathbf{Y}_{k+1} - \mathbf{X}_{k+1}^T).$$

Pseudocode for FroMax may be found in Algorithm 1. Convergence properties of FroMax are highly dependent on chosen parameters η, β_1 , and β_2 .

4. NILE-Pro. Our second algorithm, *Near-Isometric Linear Embedding via Proximal Mapping*, or *NILE-Pro* seeks to minimize the RIP constraint directly to solve for $\boldsymbol{\Psi}$. This minimization problem is solved using ADMM and a Moreau decomposition on a proximal mapping.

4.1. Optimization Framework. We formulate a new framework for NILE-Pro. We solve for our desired linear embedding $\boldsymbol{\Psi}$ directly:

$$(4.1) \quad \underset{\boldsymbol{\Psi}}{\text{minimize}} \|\mathcal{A}(\boldsymbol{\Psi}^T \boldsymbol{\Psi}) - \mathbf{1}_S\|_\infty.$$

By introducing an auxiliary variable \mathbf{q} , which simplifies our problem and allows us to apply constrained optimization techniques, we then have the following minimization

problem:

$$(4.2) \quad \begin{aligned} & \underset{\mathbf{q}, \Psi}{\text{minimize}} \quad \|\mathbf{q} - \mathbf{1}_S\|_\infty \\ & \text{subject to} \quad \mathbf{q} = \mathcal{A}(\Psi^T \Psi). \end{aligned}$$

We apply ADMM and use a Moreau decomposition on a proximal mapping to solve for updates. Like FroMax, this optimization problem is non-convex and thus, we cannot guarantee that NILE-Pro will converge to the optimal solution of (4.1). However, we demonstrate in §7 that NILE-Pro may produce stable, excellent results for synthetic and real-world data sets at a much faster rate than both FroMax and NuMax due to the simplified problem it solves.

4.2. Proximal Mapping and Moreau Decomposition. We introduce some machinery to solve this minimization problem:

DEFINITION 4.1. *The proximal mapping [17] of a closed proper convex function f is defined to be*

$$\text{prox}_f(\mathbf{x}) = \arg \min_{\mathbf{u}} (f(\mathbf{u}) + \frac{1}{2} \|\mathbf{u} - \mathbf{x}\|_2^2).$$

The proximal mapping may be interpreted as a compromise between minimizing f and staying near the original iterate x . Proximal mappings generalize projections, and are particularly useful since they often work very fast in practice.

THEOREM 4.2. *If a function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is proper, closed, and convex, then $\text{prox}_f(\mathbf{x})$ exists, well-defined, and unique for all \mathbf{x} [17].*

This theorem guarantees that these proximal mappings exist for closed, proper, and convex functions.

DEFINITION 4.3. *The convex conjugate [17] of a closed proper convex function f is defined to be*

$$f^*(\mathbf{y}) = \sup_{\mathbf{x}} (\mathbf{y}^T \mathbf{x} - f(\mathbf{x})).$$

Then for any \mathbf{x} , the following relation holds:

$$\mathbf{x} = \text{prox}_f(\mathbf{x}) + \text{prox}_{f^*}(\mathbf{x}),$$

where f^* is the convex conjugate of f . This decomposition, called the *Moreau decomposition*, generalizes the orthogonal decomposition on subspaces. We apply this machinery to help us solve for the update for \mathbf{q} .

4.3. ADMM. Following a similar method as FroMax, we relax our linear constraints and find our augmented Lagrangian of (14):

$$(4.3) \quad L_A(\Psi, \mathbf{q}; \boldsymbol{\omega}) = \|\mathbf{q} - \mathbf{1}_S\|_\infty + \frac{\beta}{2} \|\mathcal{A}(\Psi^T \Psi) - \mathbf{q} - \boldsymbol{\omega}\|_2^2.$$

Here, $\boldsymbol{\omega} \in \mathbb{R}^S$ is the scaled Lagrange multiplier. The optimization in (4.3) is carried out over $\Psi \in \mathbb{R}^{r \times N}$ and $\mathbf{q} \in \mathbb{R}^S$, while $\boldsymbol{\omega}$ is updated. Each variable is updated in an alternating fashion. The following steps below are performed until convergence.

Update Ψ : Isolating the terms that involve Ψ , we obtain a new estimate Ψ_{k+1} as the solution of the constrained optimization problem

$$(4.4) \quad \Psi_{k+1} \leftarrow \arg \min_{\Psi} \frac{\beta}{2} \|\mathcal{A}(\Psi^T \Psi) - \mathbf{q}_k - \omega_k\|_2^2.$$

Update \mathbf{q} : Isolating the terms that involve \mathbf{q} , we obtain a new estimate \mathbf{q}_{k+1} as the solution of the constrained optimization problem

$$(4.5) \quad \mathbf{q}_{k+1} \leftarrow \arg \min_{\mathbf{q}} \|\mathbf{q} - \mathbf{1}_S\|_\infty + \frac{\beta}{2} \|\mathcal{A}(\Psi_{k+1}^T \Psi_{k+1}) - \mathbf{q} - \omega_k\|_2^2.$$

Set $\mathbf{x} = \mathbf{q}_{k+1} - \mathbf{1}_S$ and $\tau = \mathcal{A}(\Psi_{k+1}^T \Psi_{k+1}) - \omega_k - \mathbf{1}_S$, let $f(\mathbf{x}) = \frac{1}{\beta} \|\mathbf{x}\|_\infty$ and its convex conjugate

$$f^*(\mathbf{x}) = \begin{cases} 0 & \text{if } \|\beta \mathbf{x}\|_1 \leq 1 \\ \infty & \text{if } \|\beta \mathbf{x}\|_1 > 1 \end{cases}.$$

Then $\text{prox}_{f^*}(\tau) = \mathcal{P}_{\{\|\mathbf{x}\|_1 \leq 1\}}(\beta \tau)$, i.e. the projection of $\beta \tau$ to the ℓ^1 ball. Then applying Moreau's decomposition, we have the \mathbf{q} update:

$$(4.6) \quad \mathbf{x} = \frac{1}{\beta} (\beta \tau - \mathcal{P}_{\{\|\mathbf{x}\|_1 \leq 1\}}(\beta \tau))$$

$$(4.7) \quad \mathbf{q} = \frac{1}{\beta} (\beta \tau - \mathcal{P}_{\{\|\mathbf{x}\|_1 \leq 1\}}(\beta \tau)) + \mathbf{1}_S.$$

Update ω : Following standard augmented Lagrangian methods, we update ω according to the following equation

$$(4.8) \quad \omega_{k+1} \leftarrow \omega_k - \beta (\mathcal{A}(\Psi_{k+1}^T \Psi_{k+1}) - \mathbf{q}_{k+1}).$$

Pseudocode for NILE-Pro may be found in Algorithm 2.

5. Rank Adjustment and Column Generation. In this section, we discuss rank adjustment and column generation heuristics. We develop rank adjustment methods to discover the lowest optimal rank for both FroMax and NILE-Pro. Abbreviating rank adjustment to *RA*, we call our rank adjusted algorithms *FroMax RA* and *NILE-Pro RA*, respectively. We also use column generation techniques following Hegde et. al. [6] to work with subsets of $S(\mathcal{X})$ to lower the memory complexity of these algorithms, which we name *FroMax CG* and *NILE-Pro CG*, respectively. We discuss each heuristic algorithm in detail below.

5.1. Rank Adjustment. Though FroMax and NILE-Pro may dramatically decrease the time of solving for projection matrix Ψ , both algorithms do not find an optimal rank for dimensionality reduction like NuMax. Hence, we propose a heuristic rank adjustment method that uses the discovered matrix $P = \Psi^T \Psi$ to give a good initialization for Ψ of lower rank.

Given a sufficiently large rank, $R_0 \gg r$, the optimal rank, we run our dimensionality reduction algorithm for a maximum number of iterations or until convergence to find Ψ . If our algorithm converges, we return $\mathbf{P} = \Psi^T \Psi$ and find $\Psi_0 = \mathbf{\Gamma}_M^{1/2} \mathbf{U}_M^T$, where $\mathbf{P} = \mathbf{U} \mathbf{\Gamma} \mathbf{U}^T$ from \mathbf{P} 's eigendecomposition. We then initialize our algorithm again with rank $R_1 = R_0 - 1$ and Ψ_0 which we found in the last iteration and test again for convergence. We continue this process until we reach the maximum number of iterations within the algorithm and return the Ψ given in the last iteration, considering its rank $r = R_k$ to be optimal. We summarize our rank adjustment heuristic in Algorithm 3.

Algorithm 2 NILE-Pro

Inputs: Secant set $\mathcal{S}(\mathcal{X}) = \{\mathbf{v}_i\}_{i=1}^S$, isometry constant δ , max iterations $m > 0$, initial rank r

Parameters: $\beta > 0$

Output: Ψ

for $k = 1, \dots, m$ **do**

$$\tau \leftarrow \mathcal{A}(\Psi_k^T \Psi_k) - \omega - \mathbf{1}_S$$

$$\mathbf{q}_{k+1} \leftarrow \frac{1}{\beta} \left(\beta \tau - \mathcal{P}_{\{\|x\|_1 \leq 1\}}(\beta \tau) \right) + \mathbf{1}_S$$

$$\Psi_{k+1} \leftarrow \Psi_k - 2\eta \Psi_k \mathcal{A}^*(\mathcal{A}(\Psi_k^T \Psi_k) - \mathbf{q}_{k+1} - \omega)$$

$$\omega_{k+1} \leftarrow \omega_k - \beta(\mathcal{A}(\Psi_{k+1}^T \Psi_{k+1}) - \mathbf{q}_{k+1})$$

$$\epsilon_0 \leftarrow \|\mathcal{A}(\Psi_{k+1}^T \Psi_{k+1}) - \mathbf{1}_S\|_\infty$$

if $\epsilon_0 < \epsilon$ **then**

break

end if

end for

return $\Psi \leftarrow \Psi_k$

Algorithm 3 FroMax/NILE-Pro RA

Inputs: Secant set $\mathcal{S}(\mathcal{X}) = \{\mathbf{v}_i\}_{i=1}^S$, isometry constant δ , max iterations for algorithm $m > 0$, initial rank R_0 , max iterations for RA M , Ψ_0

for $k = 1, \dots, M$ **do**

$$\Psi \leftarrow \text{FroMax/NILE-Pro}(S, \delta, m, R_0, \Psi_0)$$

$$P \leftarrow \Psi^T \Psi$$

$$R_{k+1} \leftarrow R_k - 1$$

$$[\Gamma, U] \leftarrow \text{eig}(P)$$

$$\Psi_0 \leftarrow \Gamma^{1/2} U^T$$

if FroMax/NILE-Pro does not converge **then**

break

end if

end for

return $\Psi, r \leftarrow R_k$

5.2. Column Generation. Since FroMax and NILE-Pro use the secants of a given data set, applications involving millions of secants may be prohibited by the memory complexity of these algorithms. Some methods that are used to address large data sets include stochastic and online methods. Stochastic methods use random subsets of the data to learn an estimate for the entire data set [9, 20]. Online methods uses sequentially available data to update the current iterate then discards the information [13, 10]. Our column generation algorithms, FroMax CG and NILE-

Algorithm 4 FroMax/NILE-Pro CG

Inputs: Secant set $\mathcal{S}(\mathcal{X}) = \{\mathbf{v}_i\}_{i=1}^S$, isometry constant δ , max iterations for algorithm $m > 0$, rank r , the FroMax or NILE-Pro algorithm

```
while not converged do
   $\hat{S} \leftarrow \{\mathbf{v}_i \in S_0 : |\mathbf{v}_i^T \Psi^T \Psi \mathbf{v}_i - 1| \geq \delta\}$ 
   $S_1 \leftarrow \{\mathbf{v}_i \in \mathcal{S}(\mathcal{X}) : \mathbf{v}_i \notin S_0\}_{i=1}^{S''}$ 
   $\hat{S} \leftarrow \hat{S} \cup \{\mathbf{v}_i \in S_1 : |\mathbf{v}_i^T \Psi^T \Psi \mathbf{v}_i - 1| \geq \delta\}$ 
   $\Psi \leftarrow \text{FroMax/NILE-Pro}(\hat{S}, \delta)$ 
   $S_0 \leftarrow \hat{S}$ 
end while

return  $\Psi$ 
```

Pro CG, combines stochastic and online methods to estimate solutions to large-scale problems.

Similar to NuMax's column generation, which is based off of the Karush-Kuhn-Tucker (KKT) conditions, we apply a simple, greedy method to rapidly find the active constraints for (3.2) or (4.1).

1. Solve (3.2) or (4.1) with a small subset $S_0 \subset \mathcal{S}(\mathcal{X})$ using FroMax (Algorithm 1) or NILE-Pro (Algorithm 2), respectively to obtain an initial estimate $\hat{\Psi}$. Identify the set \hat{S} of secants that correspond to the active constraints:

$$\hat{S} \leftarrow \{\mathbf{v}_i \in S_0 : |\mathbf{v}_i^T \hat{\Psi}^T \hat{\Psi} \mathbf{v}_i - 1| \geq \delta\}.$$

2. Select additional secants $S_1 \subset \mathcal{S}(\mathcal{X})$ not selected previously and identify all secants among S_1 that violate the constraint at the current estimate $\hat{\Psi}$. Then, append these secants to the set of active constraints \hat{S} to obtain an augmented set \hat{S}

$$\hat{S} \leftarrow \hat{S} \cup \{\mathbf{v}_i \in S_1 : |\mathbf{v}_i^T \hat{\Psi}^T \hat{\Psi} \mathbf{v}_i - 1| \geq \delta\}.$$

3. Solve (3.2) or (4.1) with the new augmented set \hat{S} using FroMax or NILE-Pro to obtain a new estimate $\hat{\Psi}$.
4. Identify the secants that correspond to active constraints and repeat Steps 2 and 3 until convergence is reached for $\hat{\Psi}$.

Column generation allows us to perform a large numerical optimization procedure on smaller subsets of $\mathcal{S}(\mathcal{X})$, resulting in significant computational gains. A key benefit of FroMax CG and NILE-Pro CG is that the subsets of secants used during each iteration never has to be explicitly stored in memory and can be generated on the fly. This leads to significant improvements in memory complexity.

However, because FroMax and NILE-Pro are already both non-convex, column generation makes these algorithms even less predictable. Though these algorithms are not guaranteed to converge to an optimal solution, they appear to yield excellent results on large, real-world data sets, as we will show in §7.

Pseudocode for FroMax/Nile-Pro CG is found in Algorithm 4. Our column generation method converges when no additional secants violate our constraint.

#DATA	δ	FroMAX		NILE-Pro	
		RANK	TIME	RANK	TIME
60	0.4	9	1.3	9	0.7
	0.25	9	1.2	9	1.1
	0.1	13	1.4	13	1.5
200	0.4	16	511.2	16	109.1
	0.25	18	269.4	18	144.4
	0.1	27	74.5	27	448.5

TABLE 7.1

Comparison of runtime performance for FroMax and NILE-Pro on $S(\mathcal{X}_1)$ and $S(\mathcal{X}_2)$ given δ and the optimal rank from Numax RA.

6. Convergence of Algorithms. Since FroMax and NILE-Pro are derived from applying the ADMM to non-convex problems, the convergence properties of these algorithms can be understood based on the known convergence properties of ADMM. For certain types of convex problems, ADMM has been shown to converge at a rate of $o(1/k)$ [3]. Recent results by Hong, et.al. and Wang, et. al. [7, 21] show that ADMM does converge for certain classes or types of non-convex problems. These results may suggest convergence of our algorithms.

7. Numerical Experiments. We demonstrate the performance of the FroMax and NILE-Pro algorithms in comparison to prior methods including NuMax. All of our experiments are performed on computers with Intel i5-650 processors and 4 GB of RAM unless otherwise specified. We test and compare the speed and accuracy of our algorithms through various tests on real-world and synthetic data sets.

7.1. Linear Low-Dimensional Embeddings. We first consider synthetic data sets \mathcal{X}_1 and \mathcal{X}_2 consisting of $7 \times 7 = 49$ and $14 \times 14 = 196$ dimensional images of translations of a white square on a black box respectively. We construct our training sets by randomly generating 60 49-dimensional images for \mathcal{X}_1 and 200 196-dimensional images for \mathcal{X}_2 . We then construct secant sets $S(\mathcal{X}_1)$ and $S(\mathcal{X}_2)$ by computing the normalized pairwise difference vectors between different images. We compare FroMax and NILE-Pro’s performance of producing linear, low-dimensional embeddings on these two data sets in Table 7.1.

Since NILE-Pro minimizes the RIP directly, NILE-Pro intuitively will converge faster for larger δ . Our experimental results match our theoretical intuition since NILE-Pro converges significantly faster for larger δ than lower δ .

FroMax experimentally converges faster for smaller δ than larger δ . Smaller δ restricts q to a smaller feasible set given by the RIP, leading to faster convergence.

Also note that both algorithms’ computational complexity scale significantly with the size of the data set due to the use of the secant set. Our runtime results comparing $S(\mathcal{X}_1)$ and $S(\mathcal{X}_2)$ reflect this.

7.2. Linear Low-Dimensional Embeddings with Rank Adjustment. In §7.1, we input a given rank for FroMax and NILE-Pro and compare their run time. However, usually the optimal rank for dimension reduction is not known, motivating the development of rank adjustment heuristics. To analyze the performance of our rank adjustment heuristic, we consider a synthetic data set \mathcal{X} comprised of $16 \times 16 = 256$ dimensional images of translations of a white square or a disk on a black box respectively, see figure 7.1. We construct a secant set $S(\mathcal{X})$ and compare PCA,

Numax RA, FroMax RA and NILE-Pro RA's performance of producing linear, low-dimensional embeddings on this data set.

Figure 7.2 plots the variation of the number of measurements M as a function of the isometry constant δ . We observe that NILE-Pro RA achieves the desired isometry constant on the secants using by far the fewest number of measurements. FroMax RA performs better for small δ due to the correlation between δ and q , as we discussed before. Moreover, both Numax RA and Fromax RA greatly outperform PCA, a popular embedding technique in the literature.




FIG. 7.1. Our synthetic training set consists of sixty 256-dimensional random generated translating disks and squares figure.




FIG. 7.2. A comparison of the isometry constant δ with the number of measurements for PCA, Numax RA, FroMax RA and NILE-Pro RA's performance of producing linear, low-dimensional embeddings.

δ	#DATA	NILE-PRO		FROMMAX		NUMAX	
		RK	TIME	RK	TIME	RK	TIME
0.4	95	7	25	9	102	12	71
	200	9	96	15	520	21	311
	500	11	710	27	2490	25	3477
0.2	95	11	28	11	111	14	56
	200	14	130	16	569	18	557
	500	18	751	40	1498	27	3517
0.1	95	15	41	15	91	16	21
	200	20	165	19	823	21	279
	500	25	1285	44	650	30	3410

TABLE 7.2

Comparison of runtime performance for FroMax RA, NILE-Pro RA, and NuMax on subsets of “5” images from MNIST.

7.3. Runtime Performance on MNIST with Rank Adjustment. In this experiment, we consider a more challenging, real-world data set, the MNIST data set, see figure 7.3. MNIST contains many digital images of handwritten digits and is a common benchmark data set for machine learning. We examine subsets of the training set for the digit “5”. We take subsets consisting of 95, 200, and 500 data points with original dimension 49.

We test runtime performance of FroMax and NILE-Pro RA on these data sets. Our results may be found in Table 2.



FIG. 7.3. Examples of 5 images from the MNIST dataset.

Our experimental results show that NILE-Pro RA may perform significantly faster and give a better optimal rank than NuMax while FroMax RA converges slower for larger data sets. This may be due to the nature of the local minima found in FroMax; the estimate for $P = \Psi^T \Psi$ given for a larger rank does not correspond to the local minima for lower ranks so that this initialization is beneficial.

Our results for FroMax RA reveal another issue with our rank adjustment method. FroMax RA appears to struggle with determining the optimal rank, sometimes performing worse than NuMax. We believe that our algorithm may be converging to local minima, which makes our rank adjustment ineffective. This issue motivates us to look into other rank adjustment methods that start at a sufficiently low rank and examine higher ranks to discover the optimal rank.

δ	RANK	NUMAX CG	TIME (HRS)	FROMAX CG	TIME (HRS)
0.4	72	3.09%	0.926	3.00%	0.411
0.25	98	3.15%	1.273	3.26%	0.811
0.1	167	3.31%	2.664	3.42%	1.358

TABLE 7.3

Comparison of misclassification rates and run-time performance of approximate nearest neighbor classifiers using NuMax CG and FroMax CG for given δ and rank on the MNIST test set.

Also, since rank adjustment for NILE-Pro is still based on the core NILE-Pro algorithm, we see that NILE-Pro RA converges in much slower time for smaller δ .

The former caveat motivates us to continue looking for better rank adjustment methods for FroMax.

7.4. Nearest Neighbor Classification on MNIST. The MNIST data set consists of 60,000 training data points and 10,000 test data points of handwritten digits [11]. The dataset contains 10 classes corresponding to each digit from 0-9. For this experiment, we use the $N = 20 \times 20 = 400$ -dimensional data set that excludes extra space at the boundaries. We use NuMax CG and FroMax CG to embed our MNIST training set into a lower dimensional space and nearest neighbor classification.

The misclassification rate of nearest neighbor classification on the unchanged data set is 3.47%. Table 7.3 gives the nearest neighbor classification misclassification rate for NuMax CG and FroMax CG for given δ and rank applied on the MNIST data set. In particular, though NuMax CG and FroMax CG give similar misclassification rates, FroMax CG has significantly better runtime performance than NuMax CG. Though a combined rank adjustment and column generation method has not been implemented for FroMax, the results suggest that FroMax may find a sufficiently good projection matrix in much less time.

7.5. Approximate Nearest Neighbors. Given a data set modeled by points in Euclidean space and a query point, *nearest neighbors* identifies the k closest points in the data set [2]. These points are usually used for further processing, such as unsupervised or supervised regression and classification.

However, as the dimension N of the data set grows, the computational cost of identifying the k nearest neighbors also becomes increasingly expensive. An alternative to computing nearest neighbors directly is to embed the data into a lower-dimensional subspace while preserving near-isometry, then applying nearest neighbor techniques. This method is called *approximate nearest neighbors*. Since NuMax, FroMax, and NILE-Pro construct low rank, near-isometric linear embeddings for a given distortion δ , they may potentially enable efficient ANN computations for high-dimensional data sets.

For this experiment, we use the LabelMe data set consisting of 4000 images of indoor and outdoor scenes [18]. We then computed GIST descriptors for each image, which are vectors of size $N = 512$ that roughly describe the overall spatial statistics of the image [15]. We then used NuMax CG and FroMax CG to estimate low rank, near-isometric linear embeddings for this data set for a given distortion parameter δ . Then we perform ANN computations on 1000 test data points in the corresponding low dimensional space. We compute embeddings of various ranks for FroMax CG to compare performance between different ranks.

Figure 7.4 demonstrates that FroMax CG generally attains similar if not better performance than NuMax CG for the same rank. In fact, our results suggest that

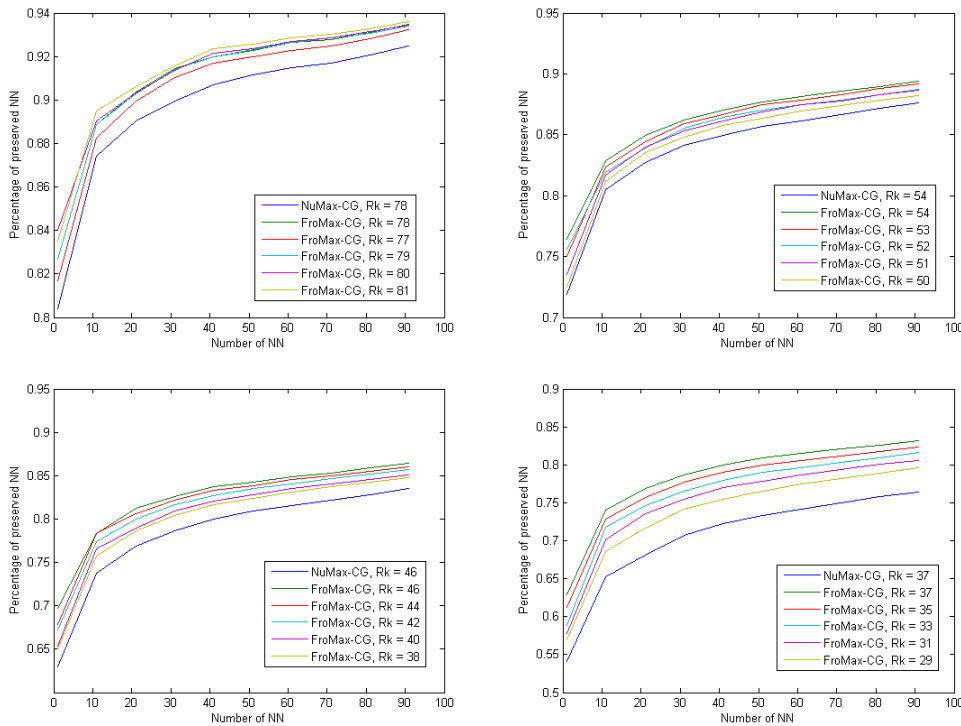


FIG. 7.4. Comparison of FroMax CG and NuMax CG on preserving nearest neighbors.

FroMax CG could perform similarly at a lower rank than NuMax CG. We leave further investigation for future research.

8. Discussion.

8.1. Research Overview. In this paper, we construct two comprehensive algorithmic frameworks for finding near-isometric linear embeddings of high-dimensional data sets. Based on the convex optimization formulation in NuMax, we proposed two non-convex minimization approaches which approximately preserve the norms of all pairwise secants of the given dataset. In particular, we developed two algorithms, FroMax and NILE-Pro, that may construct the desired embedding with smaller computational complexity than NuMax.

Since NuMax automatically discovers the optimal rank, we created a rank adjustment method for finding the best rank for our algorithms. We also implemented column generation in addition to FroMax and NILE-Pro so our algorithms can be adapted to perform on larger data sets.

Constructing linear, information-preserving embeddings of high-dimensional signals to lower-dimensional signals have become of significant importance for a wide range of machine learning and compressive sensing applications. However, little is known about near-isometric linear embeddings beyond the Johnson-Lindenstrauss Lemma. The frameworks discussed in this paper build on the convex, deterministic approach of NuMax to produce practical, potentially more computationally efficient dimension reduction algorithms that are both information-preserving and feasible for a broad range of applications. Though we do not provide an analytical foundation to

our work due to the non-convex nature of our algorithms, we hope to initiate work in developing a theoretical basis for similar work.

8.2. Future Work. There are still many challenges left to tackle. As discussed in §7, we still need to further develop rank adjustment methods for FroMax and column generation techniques for NILE-Pro. We would also like to incorporate both rank adjustment and column generation together. One direction is to consider an eigengap heuristic for rank adjustment, in which some heuristic is set based on the difference between singular values of the matrix P to determine the next chosen rank.

In addition, further testing and parameter-tweaking is necessary to analyze and optimize the stability and performance of our algorithms on various data sets. Other heuristics for non-convex optimization, such as applying perturbations to avoid local minima, may also be applied to give better solutions. We defer the study of these challenges and heuristics for future research.

Acknowledgements. The work of Jerry Luo, Kayla Shapiro, and Hao-Jun Michael Shi were supported in part by the California Research Training Program for Computational and Applied Mathematics 2014 under NSF Grant DMS-1045536. The work of Qi Yang was supported in part by USC Provost’s Undergraduate Research Fellowship and the WiSE Research Experience for Undergraduates. We would like to thank Dr. Ming Yan and Dr. Wotao Yin for their consistent support and mentorship throughout the entirety of this project.

REFERENCES

- [1] EMMANUEL J CANDES AND TERENCE TAO, *Decoding by linear programming*, Information Theory, IEEE Transactions on, 51 (2005), pp. 4203–4215.
- [2] THOMAS M COVER AND PETER E HART, *Nearest neighbor pattern classification*, Information Theory, IEEE Transactions on, 13 (1967), pp. 21–27.
- [3] WEI DENG, MING-JUN LAI, AND WOTAO YIN, *On the $o(1/k)$ convergence and parallelization of the alternating direction method of multipliers*, arXiv preprint arXiv:1312.3040, (2013).
- [4] EDWARD GRANT, CHINMAY HEGDE, AND PIOTR INDYK, *Nearly optimal linear embeddings into very low dimensions*, in Global Conference on Signal and Information Processing (GlobalSIP), 2013 IEEE, IEEE, 2013, pp. 973–976.
- [5] CHINMAY HEGDE, ASWIN C SANKARANARAYANAN, AND RICHARD G BARANIUK, *Near-isometric linear embeddings of manifolds*, in Statistical Signal Processing Workshop (SSP), 2012 IEEE, IEEE, 2012, pp. 728–731.
- [6] C. HEGDE, A. C. SANKARANARAYANAN, W. YIN, AND R. G. BARANIUK, *NuMax: a convex approach for learning near-isometric linear embeddings*, IEEE Transactions on Signal Processing, 63 (2015).
- [7] MINGYI HONG, ZHI-QUAN LUO, AND MEISAM RAZAVIYAYN, *Convergence analysis of alternating direction method of multipliers for a family of nonconvex problems*, in Acoustics, Speech and Signal Processing (ICASSP), 2015 IEEE International Conference on, IEEE, 2015, pp. 3836–3840.
- [8] WILLIAM JOHNSON AND JORAM LINDENSTRAUSS, *Extensions of Lipschitz mappings into a Hilbert space*, in Conference in modern analysis and probability (New Haven, Conn., 1982), vol. 26 of Contemporary Mathematics, American Mathematical Society, 1984, pp. 189–206.
- [9] HAROLD KUSHNER AND G GEORGE YIN, *Stochastic approximation and recursive algorithms and applications*, vol. 35, Springer Science & Business Media, 2003.
- [10] LEON BOTTOU YANN LE CUN AND LEON BOTTOU, *Large scale online learning*, Advances in neural information processing systems, 16 (2004), p. 217.
- [11] YANN LECUN, CORINNA CORTES, AND CHRISTOPHER JC BURGESS, *The MNIST database of handwritten digits*, 1998.
- [12] DANIEL D. LEE AND H. SEBASTIAN SEUNG, *Algorithms for non-negative matrix factorization*, in Advances in Neural Information Processing Systems 13, T.K. Leen, T.G. Dietterich, and V. Tresp, eds., MIT Press, 2001, pp. 556–562.

- [13] JULIEN MAIRAL, FRANCIS BACH, JEAN PONCE, AND GUILLERMO SAPIRO, *Online learning for matrix factorization and sparse coding*, J. Mach. Learn. Res., 11 (2010), pp. 19–60.
- [14] BRUCE MOORE, *Principal component analysis in linear systems: Controllability, observability, and model reduction*, Automatic Control, IEEE Transactions on, 26 (1981), pp. 17–32.
- [15] AUDE OLIVA AND ANTONIO TORRALBA, *Modeling the shape of the scene: A holistic representation of the spatial envelope*, International journal of computer vision, 42 (2001), pp. 145–175.
- [16] BENJAMIN RECHT, MARYAM FAZEL, AND PABLO A PARRILO, *Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization*, SIAM review, 52 (2010), pp. 471–501.
- [17] R. TYRRELL ROCKAFELLAR, *Convex analysis*, 1970.
- [18] BRYAN C RUSSELL, ANTONIO TORRALBA, KEVIN P MURPHY, AND WILLIAM T FREEMAN, *Labelme: a database and web-based tool for image annotation*, International journal of computer vision, 77 (2008), pp. 157–173.
- [19] ALI SADEGHIAN, BUBACARR BAH, AND VOLKAN CEVHER, *Energy-aware adaptive bi-lipschitz embeddings*, in 10th International Conference on Sampling Theory and Applications (SampTA), 2013.
- [20] JAMES C SPALL, *Introduction to stochastic search and optimization: estimation, simulation, and control*, vol. 65, John Wiley & Sons, 2005.
- [21] YU WANG, WOTAO YIN, AND JINSHAN ZENG, *Global convergence of admm in nonconvex nonsmooth optimization*, arXiv preprint arXiv:1511.06324, (2015).
- [22] YANGYANG XU, WOTAO YIN, ZAIWEN WEN, AND YIN ZHANG, *An alternating direction algorithm for matrix completion with nonnegative factors*, Frontiers of Mathematics in China, 7 (2012), pp. 365–384.