

A Majorization-Minimization Gauss-Newton Method for 1-Bit Matrix Completion

Xiaoqian Liu, Xu Han, Eric C. Chi & Boaz Nadler

To cite this article: Xiaoqian Liu, Xu Han, Eric C. Chi & Boaz Nadler (09 Jan 2025): A Majorization-Minimization Gauss-Newton Method for 1-Bit Matrix Completion, Journal of Computational and Graphical Statistics, DOI: [10.1080/10618600.2024.2428610](https://doi.org/10.1080/10618600.2024.2428610)

To link to this article: <https://doi.org/10.1080/10618600.2024.2428610>



© 2025 The Author(s). Published with license by Taylor & Francis Group, LLC.



[View supplementary material](#)



Published online: 09 Jan 2025.



[Submit your article to this journal](#)



Article views: 368



[View related articles](#)



[View Crossmark data](#)

A Majorization-Minimization Gauss-Newton Method for 1-Bit Matrix Completion

Xiaoqian Liu^a, Xu Han^b, Eric C. Chi^c, and Boaz Nadler^d

^aDepartment of Statistics, University of California, Riverside, Riverside, CA; ^bTencent, Beijing, China; ^cDepartment of Statistics, Rice University, Houston, TX; ^dDepartment of Computer Science and Applied Mathematics, Weizmann Institute of Science, Rehovot, Israel

ABSTRACT

In 1-bit matrix completion, the aim is to estimate an underlying low-rank matrix from a partial set of binary observations. We propose a novel method for 1-bit matrix completion called Majorization-Minimization Gauss-Newton (MMGN). Our method is based on the majorization-minimization principle, which converts the original optimization problem into a sequence of standard low-rank matrix completion problems. We solve each of these sub-problems by a factorization approach that explicitly enforces the assumed low-rank structure and then apply a Gauss-Newton method. Using simulations and a real data example, we illustrate that in comparison to existing 1-bit matrix completion methods, MMGN outputs comparable if not more accurate estimates. In addition, it is often significantly faster, and less sensitive to the spikiness of the underlying matrix. In comparison with three standard generic optimization approaches that directly minimize the original objective, MMGN also exhibits a clear computational advantage, especially when the fraction of observed entries is small. Supplementary materials for this article are available online.

ARTICLE HISTORY

Received April 2023
Accepted October 2024

KEYWORDS

Binary observations;
Constrained least squares;
Low-rank matrix; Maximum
likelihood estimate

1. Introduction

In this work, we consider the following 1-bit matrix completion problem: Let $\Theta^* \in \mathbb{R}^{m \times n}$ be an unknown matrix of exact rank r^* . The observed data are a subset $\Omega \subset [m] \times [n]$ and a binary matrix \mathbf{Y} . At entries $(i, j) \notin \Omega$, $y_{ij} = 0$. At entries $(i, j) \in \Omega$, the observations are independent draws as follows,

$$y_{ij} = \begin{cases} +1 & \text{with probability } \Phi(\theta_{ij}^*), \\ -1 & \text{with probability } 1 - \Phi(\theta_{ij}^*), \end{cases} \quad \text{for } (i, j) \in \Omega \quad (1)$$

where θ_{ij}^* is the ij th entry of Θ^* , and $\Phi: \mathbb{R} \rightarrow [0, 1]$ is a known cumulative distribution function (CDF). Given the function Φ and the matrix \mathbf{Y} , the goal of 1-bit matrix completion is to estimate the underlying matrix Θ^* , see for example Davenport et al. (2014).

This problem is a variant of the standard matrix completion problem, in which one directly observes the entries θ_{ij}^* for $(i, j) \in \Omega$. In some applications, one can only observe a stochastic value y_{ij} that is non-linearly related to the underlying value θ_{ij}^* . For example, in the Netflix challenge (Koren, Bell, and Volinsky 2009), the goal was to predict the missing entries of a rating matrix whose observed entries were integers ranging from one to five. In the simplest model for recommendation systems, user responses take on only one of two values: “like” or “dislike”, leading to a 1-bit matrix completion problem. Additional applications involving completion of partially observed binary matrices include quantum state tomography (Gross et al. 2010) and link prediction in network analysis (Liben-Nowell and Kleinberg 2003).

Given the partially observed binary matrix \mathbf{Y} and an assumed target rank r , perhaps the most natural approach to estimate Θ^* is by maximizing the likelihood function, or equivalently minimizing the negative log-likelihood $\ell(\Theta)$ under a rank- r constraint,

$$\underset{\Theta \in \mathbb{R}^{m \times n}}{\text{minimize}} \ell(\Theta) \quad \text{subject to} \quad \text{rank}(\Theta) \leq r. \quad (2)$$



Specifically, under the generative model (1), $\ell(\Theta)$ is given by

$$\ell(\Theta) = - \sum_{(i,j) \in \Omega} \left\{ \delta_{ij} \log \Phi(\theta_{ij}) + (1 - \delta_{ij}) \log [1 - \Phi(\theta_{ij})] \right\}, \quad (3)$$

where $\delta_{ij} = \frac{1}{2}(1 + y_{ij})$ are rescalings of the observed entries.

1.1. Prior Work

Several authors derived algorithms as well as theoretical guarantees for the above 1-bit matrix completion problem. To ensure the 1-bit matrix completion problem is well-posed, assumptions are needed on the underlying matrix Θ^* and on the subset Ω of observed entries. First, the entries of Θ^* cannot be too large in magnitude, so that the observed entries provide sufficient information on the underlying values of Θ^* . Second, the rank of Θ^* , or a proxy thereof, also cannot be too large. Finally, the matrix needs to be de-localized, and cannot be too spiky. Here, we adopt the notion of matrix spikiness from Negahban and Wainwright (2012). For a nonzero matrix $\Theta \in \mathbb{R}^{m \times n}$, its

CONTACT Xiaoqian Liu  xiaoqian.liu@ucr.edu  Department of Statistics, University of California, Riverside, Riverside, CA 92521.

 Supplementary materials for this article are available online. Please go to www.tandfonline.com/r/JCGS.

© 2025 The Author(s). Published with license by Taylor & Francis Group, LLC.

This is an Open Access article distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivatives License (<http://creativecommons.org/licenses/by-nc-nd/4.0/>), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the original work is properly cited, and is not altered, transformed, or built upon in any way. The terms on which this article has been published allow the posting of the Accepted Manuscript in a repository by the author(s) or with their consent.

spikiness ratio $s(\Theta)$ is defined as $s(\Theta) = \sqrt{mn}\|\Theta\|_\infty/\|\Theta\|_F$ where $\|\Theta\|_\infty = \max_{ij}|\theta_{ij}|$ and $\|\Theta\|_F$ is its Frobenius norm. By definition, the spikiness ratio satisfies $1 \leq s(\Theta) \leq \sqrt{mn}$. In general, a lower value of $s(\Theta^*)$ yields an easier 1-bit matrix completion problem.

Additionally, the set Ω of observed entries needs to be well spread-out. Several previous works assumed that Ω is distributed uniformly at random. Then if $|\Omega|$ is sufficiently large and the above assumptions on Θ^* hold, it is possible to accurately estimate Θ^* , see for example Bhaskar and Javanmard (2015), Davenport et al. (2014), and Cai and Zhou (2013).

Instead of the original problem (2), some works minimized (3) subject to a constraint on the rank and an additional constraint on $\|\Theta\|_\infty$ (Bhaskar and Javanmard 2015; Ni and Gu 2016). Other works employed a constraint on a convex relaxation of the rank in conjunction with the constraint on $\|\Theta\|_\infty$. For example, Davenport et al. (2014) replaced the rank by the trace norm $\|\Theta\|_*$. They minimized the negative log-likelihood (3) under two constraints $\|\Theta\|_* \leq \alpha\beta\sqrt{mn}$ and $\|\Theta\|_\infty \leq \alpha$, where α and β are two tuning parameters. Instead of the trace norm, Cai and Zhou (2013) employed the max-norm (Linial et al. 2007) as a convex relaxation of the rank. They minimized the objective (3) under the two constraints $\|\Theta\|_{\max} \leq \alpha\beta$ and $\|\Theta\|_\infty \leq \alpha$, where $\|\Theta\|_{\max} = \min_{\Theta=UV^T} \{\|\mathbf{U}\|_{2,\infty}, \|\mathbf{V}\|_{2,\infty}\}$ and $\|\mathbf{U}\|_{2,\infty}$ is the largest l_2 -norm of the rows in \mathbf{U} .

Most algorithms for these prior works, however, are computationally intensive and nontrivial to implement. Bhaskar and Javanmard (2015) employed a log-barrier penalty to approximate the infinity norm constraint in combination with a cross-validation procedure to determine how to scale the penalty. The algorithm in Davenport et al. (2014), at each iteration, runs an inner ADMM algorithm to compute the projection onto the intersection of an infinity norm ball and nuclear norm ball. The algorithm in Cai and Zhou (2013), at each iteration, requires projection onto a max-norm constraint which involves solving a nonconvex optimization problem iteratively in conjunction with a rank selection step. Hence, some of these methods do not scale well to large matrices and are prohibitively slow. As illustrated later on in Section 4, publicly available implementations may take hours to complete medium-sized matrices.

1.2. Our Contributions

The main contribution of this work is the derivation of a simple and computationally fast approach to minimize (2). Our iterative method, denoted Majorization-Minimization Gauss-Newton (MMGN), relies on a majorization-minimization strategy. This leads to solving a sequence of standard simpler matrix completion problems. We approximately solve these problems by a single Gauss-Newton step, each of which requires solving a least squares problem. Hence, our method is simple to implement, requiring only standard linear algebra tools. MMGN is typically at least an order of magnitude faster than existing 1-bit matrix completion methods while achieving comparable or even superior estimation accuracy. In addition, MMGN is faster than generic optimization methods that minimize the original objective (2), including gradient descent, quasi-Newton methods, and manifold optimization algorithms. Beyond its computational

advantages, as illustrated via simulations, MMGN's performance degrades more slowly as the spikiness of the underlying matrix increases compared to alternative methods.

1.3. Notation and Organization

We denote vectors by boldface lowercase letters and matrices by boldface capital letters, for example, $\mathbf{a} \in \mathbb{R}^n$ and $\mathbf{A} \in \mathbb{R}^{m \times n}$. We denote the entries of a vector \mathbf{a} and matrix \mathbf{A} by a_i and a_{ij} , respectively. The rank of a matrix \mathbf{A} is denoted by $\text{rank}(\mathbf{A})$, its Moore-Penrose pseudo inverse by \mathbf{A}^\dagger , and its column-major vectorization, that is, the vector obtained by stacking the columns of \mathbf{A} one after the other, by $\text{vec}(\mathbf{A})$. We will use the following seminorm of a matrix \mathbf{A} , denoted by $\|\mathbf{A}\|_{F(\Omega)} = \sqrt{\sum_{(i,j) \in \Omega} a_{ij}^2}$, where Ω is a subset of \mathbf{A} 's indices. Similarly, for a vector \mathbf{a} , we denote $\|\mathbf{a}\|_\Omega = \sqrt{\sum_{i \in \Omega} a_i^2}$, where Ω is a subset of \mathbf{a} 's indices. For two matrices \mathbf{A} and \mathbf{B} of the same size, we denote the Hadamard product by $\mathbf{A} \circ \mathbf{B}$, namely $(\mathbf{A} \circ \mathbf{B})_{ij} = a_{ij}b_{ij}$. Similarly, the element-wise quotient of \mathbf{A} and \mathbf{B} is $\mathbf{A} \oslash \mathbf{B}$, namely $(\mathbf{A} \oslash \mathbf{B})_{ij} = a_{ij}/b_{ij}$. For a set S , we denote its cardinality by $|S|$. Finally, for any univariate function $\Phi: \mathbb{R} \rightarrow \mathbb{R}$, we denote the matrix obtained by applying Φ entry-wise to \mathbf{A} by $\Phi(\mathbf{A})$.

The rest of the article is organized as follows. Section 2 introduces the MMGN method. Section 3 discusses several related problems and approaches similar to ours to solve them. Section 4 presents an empirical evaluation of MMGN with simulations and a real data example. Section 5 concludes with a discussion.

2. The MMGN Method

We first present an overview of our approach. Given a target rank r , our strategy for solving the nonconvex problem (2) is based on the majorization-minimization (MM) principle. Concretely, we replace the objective $\ell(\Theta)$ with a simpler surrogate function, called a majorization. As we show below, minimizing our majorization leads to a standard low-rank matrix completion problem. Hence, MMGN solves a sequence of optimization problems of the form

$$\underset{\Theta \in \mathbb{R}^{m \times n}}{\text{minimize}} \quad \|\Theta - \tilde{\mathbf{Y}}\|_{F(\Omega)}^2 \quad \text{subject to} \quad \text{rank}(\Theta) \leq r, \quad (4)$$

where the matrix $\tilde{\mathbf{Y}}$ depends on the observed matrix \mathbf{Y} and on the current estimate $\tilde{\Theta}$.

We solve (4) by a factorization approach, expressing $\Theta = \mathbf{UV}^T$ where $\mathbf{U} \in \mathbb{R}^{m \times r}$ and $\mathbf{V} \in \mathbb{R}^{n \times r}$. Therefore, we rewrite (4) as the following equivalent unconstrained problem

$$\underset{\mathbf{U} \in \mathbb{R}^{m \times r}, \mathbf{V} \in \mathbb{R}^{n \times r}}{\text{minimize}} \quad \|\mathbf{UV}^T - \tilde{\mathbf{Y}}\|_{F(\Omega)}^2. \quad (5)$$

Finally, we solve (5) using a modified Gauss-Newton algorithm introduced by Zilber and Nadler (2022).

2.1. The Majorization-Minimization Principle

The MM principle (De Leeuw 1994; Heiser 1995; Lange, Hunter, and Yang 2000; Hunter and Lange 2004; Lange 2016) converts the minimization of a challenging function $\ell(\Theta)$ into a sequence

of simpler optimization problems. We approximate $\ell(\Theta)$ by a surrogate function or majorization $g(\Theta | \tilde{\Theta})$ anchored at the current estimate $\tilde{\Theta}$. The majorization $g(\Theta | \tilde{\Theta})$ needs to satisfy two conditions: (i) a tangency condition $g(\tilde{\Theta} | \tilde{\Theta}) = \ell(\tilde{\Theta})$ for all $\tilde{\Theta}$ and (ii) a domination condition $g(\Theta | \tilde{\Theta}) \geq \ell(\Theta)$ for all Θ . The associated MM algorithm is defined by the iterates

$$\Theta_{t+1} = \arg \min_{\Theta} g(\Theta | \Theta_t), \quad t = 0, 1, \dots \quad (6)$$

The tangency and domination conditions imply that

$$\ell(\Theta_{t+1}) \leq g(\Theta_{t+1} | \Theta_t) \leq g(\Theta_t | \Theta_t) = \ell(\Theta_t).$$

In other words, the objective function values of the MM iterates decrease monotonically. Finding the global minimizer of $g(\Theta | \Theta_t)$ is not necessary to ensure the descent property. Therefore, one can inexactly solve (6) and still guarantee a monotonic decrease of the objective. A key component for the success of the MM principle is the construction of a surrogate function that is easy to minimize.

Recall that we seek to minimize the negative log-likelihood of the 1-bit matrix completion problem (3), which we restate here for convenience:

$$\ell(\Theta) = - \sum_{(i,j) \in \Omega} \left\{ \delta_{ij} \log \Phi(\theta_{ij}) + (1 - \delta_{ij}) \log [1 - \Phi(\theta_{ij})] \right\},$$

where $\delta_{ij} = \frac{1}{2}(1 + y_{ij}) \in \{0, 1\}$. To derive our majorization, we assume the CDF $\Phi(\theta)$ in (3) satisfies the following two conditions.

- A1.** The function $\log \Phi(\theta)$ is L -Lipschitz differentiable.
- A2.** The density function $\phi(\theta) = \Phi'(\theta)$ is symmetric around zero. This implies that $\Phi(\theta) = 1 - \Phi(-\theta)$.

The following proposition describes a quadratic majorization under the above assumptions.

Proposition 2.1. Let $\Phi(\theta)$ be a CDF that satisfies assumptions A1 and A2. The following is a majorization of $\ell(\Theta)$ at $\tilde{\Theta}$

$$g(\Theta | \tilde{\Theta}) = \frac{L}{2} \|\Theta - \tilde{\mathbf{Y}}\|_{\mathbb{F}(\Omega)}^2 + c(\tilde{\Theta}), \quad (7)$$

where $c(\tilde{\Theta})$ depends on $\tilde{\Theta}$ but not on Θ , and

$$\tilde{\mathbf{Y}} = \tilde{\Theta} + \frac{1}{L} \left(\mathbf{Y} \circ \phi(\tilde{\Theta}) \right) \oslash \Phi \left(\mathbf{Y} \circ \tilde{\Theta} \right). \quad (8)$$

Two popular CDFs in 1-bit matrix completion are $\Phi(x) = 1/(1 + e^{-x/\sigma})$ and $\Phi(x) = \frac{1}{\sqrt{2\pi\sigma^2}} \int_{-\infty}^x \exp\left(-\frac{w^2}{2\sigma^2}\right) dw$, which correspond to logistic and Gaussian random variables, respectively. Both satisfy assumptions A1 and A2. Hence, by [Proposition 2.1](#) we obtain the following specific quadratic majorizations.

Corollary 2.1. The logistic model with $\Phi(x) = 1/(1 + e^{-x/\sigma})$ satisfies assumptions A1 and A2 with Lipschitz constant $L = \frac{1}{4\sigma^2}$. Hence, a majorization is given by (7) with

$$\tilde{\mathbf{Y}} = \tilde{\Theta} + 4\sigma \mathbf{Y} \circ \Phi \left(-\mathbf{Y} \circ \tilde{\Theta} \right).$$

Corollary 2.2. The probit model with $\Phi(x) = \frac{1}{\sqrt{2\pi\sigma^2}} \int_{-\infty}^x \exp\left(-\frac{w^2}{2\sigma^2}\right) dw$ satisfies assumptions A1 and A2 with Lipschitz constant $L = \frac{1}{\sigma^2}$. Hence, a majorization is given by (7) with

$$\tilde{\mathbf{Y}} = \tilde{\Theta} + \sigma^2 \mathbf{Y} \circ \phi(\tilde{\Theta}) \oslash \Phi(\mathbf{Y} \circ \tilde{\Theta}).$$

Proofs of [Proposition 2.1](#), [Corollaries 2.1](#), and [2.2](#) are in the supplementary materials. Combined with the low-rank constraint, these corollaries indicate that both the logistic and probit models lead to MM updates that require solving a standard low-rank matrix completion problem.

$$\underset{\Theta \in \mathbb{R}^{m \times n}}{\text{minimize}} \quad \|\Theta - \tilde{\mathbf{Y}}\|_{\mathbb{F}(\Omega)}^2 \quad \text{subject to} \quad \text{rank}(\Theta) \leq r. \quad (9)$$

We denote the solution to the problem (9) by Θ_{MM} . We next review how we approximately solve the non-convex problem (9) by a Gauss-Newton strategy.

2.2. Inexact Majorization-Minimization via Gauss-Newton

Let $g(\Theta | \tilde{\Theta})$ be a quadratic majorization of the form (7) with $\tilde{\mathbf{Y}}$ defined in (8), where $\tilde{\Theta}$ is the current estimate of Θ . It is useful to express $g(\Theta | \tilde{\Theta})$ in terms of the difference $\Theta - \tilde{\Theta}$. The MM-update in (9) can then be written as

$$\begin{aligned} \Theta_{\text{MM}} &= \arg \min_{\text{rank}(\Theta) \leq r} g(\Theta | \tilde{\Theta}) \\ &= \arg \min_{\text{rank}(\Theta) \leq r} \left\| \Theta - \tilde{\Theta} - \mathbf{X} \right\|_{\mathbb{F}(\Omega)}^2, \end{aligned} \quad (10)$$

where

$$\mathbf{X} = \tilde{\mathbf{Y}} - \tilde{\Theta} = \frac{1}{L} \mathbf{Y} \circ \phi(\tilde{\Theta}) \oslash \Phi(\mathbf{Y} \circ \tilde{\Theta}). \quad (11)$$

We solve (10) by the following factorization approach. We express $\tilde{\Theta}$ and Θ as the products of two rank- r factor matrices with

$$\tilde{\Theta} = \tilde{\mathbf{U}}\tilde{\mathbf{V}}^T \quad \text{and} \quad \Theta = (\tilde{\mathbf{U}} + \Delta\mathbf{U})(\tilde{\mathbf{V}} + \Delta\mathbf{V})^T.$$

The variables $\tilde{\mathbf{U}} \in \mathbb{R}^{m \times r}$ and $\tilde{\mathbf{V}} \in \mathbb{R}^{n \times r}$ are the factor matrices corresponding to the current estimate $\tilde{\Theta}$, whereas $\Delta\mathbf{U} \in \mathbb{R}^{m \times r}$ and $\Delta\mathbf{V} \in \mathbb{R}^{n \times r}$ are their updates. Therefore, problem (10) can be written in terms of the new variables $(\Delta\mathbf{U}, \Delta\mathbf{V})$ as

$$\underset{\Delta\mathbf{U}, \Delta\mathbf{V}}{\text{minimize}} \quad \left\| \mathbf{X} - \tilde{\mathbf{U}}\Delta\mathbf{V}^T - \Delta\mathbf{U}\tilde{\mathbf{V}}^T - \Delta\mathbf{U}\Delta\mathbf{V}^T \right\|_{\mathbb{F}(\Omega)}^2. \quad (12)$$

The optimization problem in (12) is a nonlinear least squares problem. Motivated by [Zilber and Nadler \(2022\)](#), we neglect the second order term $\Delta\mathbf{U}\Delta\mathbf{V}^T$. We thus compute the solution $(\Delta\mathbf{U}^*, \Delta\mathbf{V}^*)$ to

$$\underset{\Delta\mathbf{U}, \Delta\mathbf{V}}{\text{minimize}} \quad \left\| \mathbf{X} - \tilde{\mathbf{U}}\Delta\mathbf{V}^T - \Delta\mathbf{U}\tilde{\mathbf{V}}^T \right\|_{\mathbb{F}(\Omega)}^2. \quad (13)$$

Approximating the nonlinear least squares problem with the linear least squares problem in (13) corresponds to taking a single Gauss-Newton iteration. This approach produces the following inexact solution of (10), denoted $\hat{\Theta}_{\text{MM}}$,

$$\hat{\Theta}_{\text{MM}} = (\tilde{\mathbf{U}} + \Delta\mathbf{U}^*)(\tilde{\mathbf{V}} + \Delta\mathbf{V}^*)^T. \quad (14)$$

Algorithm 1 MMGN**Input:** Ω, \mathbf{Y} , target rank r , tolerance tol

- 1: Initialize $\mathbf{U}_0, \mathbf{V}_0$, set $t = 0$, $\text{rel} = \text{INF}$
- 2: **while** $\text{rel} > \text{tol}$ **do**
- 3: Construct \mathbf{X}_t in (11) as

▷ Majorization

$$\mathbf{X}_t = \frac{1}{L} \mathbf{Y} \circ \phi(\mathbf{U}_t \mathbf{V}_t^\top) \circ \Phi(\mathbf{Y} \circ (\mathbf{U}_t \mathbf{V}_t^\top)).$$

- 4: Compute the least l_2 -norm solution $(\Delta \mathbf{U}^*, \Delta \mathbf{V}^*)$ to

▷ Single MM-gradient step

$$\underset{\Delta \mathbf{U}, \Delta \mathbf{V}}{\text{minimize}} \quad \|\mathbf{X}_t - \mathbf{U}_t \Delta \mathbf{V}^\top - \Delta \mathbf{U} \mathbf{V}_t^\top\|_{\text{F}(\Omega)}^2.$$

- 5: Update factor matrices

$$\begin{bmatrix} \mathbf{U}_{t+1} \\ \mathbf{V}_{t+1} \end{bmatrix} = \begin{bmatrix} \mathbf{U}_t \\ \mathbf{V}_t \end{bmatrix} + \alpha_t \begin{bmatrix} \Delta \mathbf{U}^* \\ \Delta \mathbf{V}^* \end{bmatrix},$$

where α_t is the stepsize chosen by Armijo backtracking line search.

- 6: Compute the relative change in the log-likelihood

$$\text{rel} = \frac{|\ell(\mathbf{U}_{t+1} \mathbf{V}_{t+1}^\top) - \ell(\mathbf{U}_t \mathbf{V}_t^\top)|}{|\ell(\mathbf{U}_t \mathbf{V}_t^\top)|}.$$

- 7: $t = t + 1$.
- 8: **end while**
- 9: **return** $\hat{\Theta} = \mathbf{U}_t \mathbf{V}_t^\top$.

We address three important details about our approach. First, the linear least squares problem (13) has infinitely many solutions. For example, suppose that $(\Delta \mathbf{U}^*, \Delta \mathbf{V}^*)$ is a solution to (13), then $(\Delta \mathbf{U}^* + \tilde{\mathbf{U}}\mathbf{R}, \Delta \mathbf{V}^* - \tilde{\mathbf{V}}\mathbf{R}^\top)$ is also a solution for any $\mathbf{R} \in \mathbb{R}^{r \times r}$. Here we adopt the strategy used by Bauch, Nadler, and Zilber (2021) and Zilber and Nadler (2022) and select the least l_2 -norm solution, namely $(\Delta \mathbf{U}^*, \Delta \mathbf{V}^*)$ with smallest value of $\|\Delta \mathbf{U}^*\|_{\text{F}}^2 + \|\Delta \mathbf{V}^*\|_{\text{F}}^2$ among all $(\Delta \mathbf{U}^*, \Delta \mathbf{V}^*)$ that solve (13). This solution can be computed efficiently using the LSQR algorithm (Paige and Saunders 1982) or the conjugate gradient method (Hestenes and Stiefel 1952; Kammerer and Nashed 1972) applied to the normal equations.

The second detail is that the update in (14) may not necessarily lead to a decrease in the objective function. The Gauss-Newton method is an instance of the steepest descent algorithm, and consequently the solution $(\Delta \mathbf{U}^*, \Delta \mathbf{V}^*)$ corresponds to a descent direction. In fact, we prove in the supplementary materials that any solution to (13) is a descent direction. The update in (14), however, corresponds to a full Gauss-Newton step which may not necessarily decrease the value of the original objective function in (10). To overcome this potential problem, if $\hat{\Theta}_{\text{MM}}$ in (14) does not decrease the original objective, we apply the Armijo backtracking line search to select a suitable stepsize (Nocedal and Wright 2006, chap. 3). The Armijo backtracking line search requires computing the objective value, which costs $\mathcal{O}(|\Omega|r)$ flops. In our experience, however, the backtracking line search is seldom triggered in practice. Details on this and a discussion on the convergence of MMGN are in the supplementary materials.

The third detail is that there is no need to globally minimize the majorization (10). Motivated by computational con-

siderations, we inexactly solve the MM optimization problem by taking a *single* Gauss-Newton step. This differs from the approach of Zilber and Nadler (2022), which solves (13) using multiple iteration steps and is thus more computationally intensive. Inexact minimization is a standard approach in cases where exact minimization requires an iterative solver. See, for example, the split-feasibility algorithm in Xu et al. (2018), which also employed a single Gauss-Newton step. Under suitable regularity conditions, little is lost by a single-step MM-gradient approach. Not only does it avoid potentially expensive inner iterations within outer MM iterations, but it often exhibits the same local convergence rate as exact minimization (Hunter and Lange 2004).

We present the complete MMGN algorithm in Algorithm 1. MMGN requires $\mathcal{O}(|\Omega|r)$ operations per iteration, which is comparable to the gradient-based methods. Details on these per-iteration costs can be found in Section C of the supplement. Empirically, however, MMGN consistently outperforms generic gradient-based methods in terms of runtime as shown in Section 4. We close this section with a brief discussion to understand why this might be expected.

For simplicity, we work with a bivariate instance of the objective function for the factorization approach. Let $\ell : \mathbb{R} \mapsto \mathbb{R}$ be twice differentiable with an L -Lipschitz first derivative. Define $f : \mathbb{R} \times \mathbb{R} \mapsto \mathbb{R}$ as $f(u, v) = \ell(uv)$. Let $u = \tilde{u} + \Delta u$ and $v = \tilde{v} + \Delta v$ and consider the second order Taylor expansion of f around (\tilde{u}, \tilde{v}) .

$$\begin{aligned} f(u, v) &= \ell(\tilde{u}\tilde{v}) + \ell'(\tilde{u}\tilde{v})[\Delta u\tilde{v} + \tilde{u}\Delta v] \\ &\quad + \frac{\ell''(\tilde{u}\tilde{v})}{2}[\Delta u\tilde{v} + \tilde{u}\Delta v]^2 + r(\Delta u, \Delta v), \end{aligned} \quad (15)$$

where $r(\Delta u, \Delta v)$ is $o((\Delta u + \Delta v)^3)$.

MMGN employs the following majorization of $f(u, v)$ at (\tilde{u}, \tilde{v})

$$g(u, v | \tilde{u}, \tilde{v}) = \ell(\tilde{u}\tilde{v}) + \ell'(\tilde{u}\tilde{v})[\Delta u\tilde{v} + \tilde{u}\Delta v] + \frac{L}{2}[\Delta u\tilde{v} + \tilde{u}\Delta v]^2. \quad (16)$$

There are two differences between (15) and (16). First, MMGN uses the uniform bound L in lieu of computing the second derivative $\ell''(\tilde{u}\tilde{v})$. This substitution greatly simplifies computations in the general case when ℓ is a function of matrices $\mathbf{U} \in \mathbb{R}^{m \times r}$ and $\mathbf{V} \in \mathbb{R}^{n \times r}$. Second, MMGN drops the third order error term $r(\Delta u, \Delta v)$. Recall that the minimizer to (16) is not unique and that MMGN chooses the least 2-norm minimizer. This has the beneficial side-effect of finding a descent direction along which the error term $r(\Delta u, \Delta v)$ is small. In other words, MMGN picks a descent direction, among many possible descent directions, along which the surrogate function has a small approximation error. This small error is likely the reason the full Gauss-Newton step is almost always taken by MMGN.

By contrast, the gradient descent surrogate is

$$h(u, v | \tilde{u}, \tilde{v}) = \ell(\tilde{u}\tilde{v}) + \ell'(\tilde{u}\tilde{v})[\Delta u\tilde{v} + \tilde{u}\Delta v] + \frac{1}{2\alpha}[\Delta u^2 + \Delta v^2], \quad (17)$$

where α is the step-size parameter.

There are again two differences to note. First, (17) uses the isotropic quadratic function

$$\frac{1}{2\alpha}(\Delta u^2 + \Delta v^2)$$

to approximate the anisotropic quadratic function

$$\frac{\ell''(\tilde{u}\tilde{v})}{2}(\Delta u\tilde{v} + \tilde{u}\Delta v)^2$$

in (15). Moreover, the quadratic term in (17) does not even depend on the current iterate (\tilde{u}, \tilde{v}) . By contrast, MMGN uses a similar anisotropic quadratic function in (16) as in (15). The quadratic dependency in (16) on (\tilde{u}, \tilde{v}) and $(\Delta u, \Delta v)$ is the same up to a scalar multiple as in (15). Second, the update direction $(\Delta u, \Delta v)$ chosen by gradient descent can be large when far from a solution. Thus, $r(\Delta u, \Delta v)$ could be large and so the surrogate approximation in (17) may be poor along the gradient direction. This is a well-known issue with gradient descent that typically results in gradient descent taking very small steps when the objective function has steep and narrow valleys.

Now recall that the per-iteration cost of MMGN is $O(|\Omega|r)$. The cost of computing the gradient of $\ell(\mathbf{U}\mathbf{V}^T)$ is also $O(|\Omega|r)$. The above simple analysis suggests that MMGN computes updates from surrogate models that capture useful second order information with a per-iteration cost that is roughly the same cost as a gradient evaluation. In other words, we expect MMGN to behave as a quasi-Newton method that has a per-iteration cost similar to a first order method.

3. Related Problems

The 1-bit matrix completion problem as well as our approach to solving it are related to several other problems, as we outline below.

Logistic PCA. In logistic PCA, the entries of a binary data matrix $\mathbf{Y} \in \{-1, 1\}^{m \times n}$ are modeled as independent draws $y_{ij} \sim \text{Bernoulli}(\Phi(\theta_{ij}^*))$ where $\Theta^* \in \mathbb{R}^{m \times n}$ is rank- r (Collins, Dasgupta, and Schapire 2001; Schein, Saul, and Ungar 2003). Logistic PCA shares the same modeling assumptions with 1-bit matrix completion but differs in two ways. First, as its name suggests, in 1-bit matrix completion, only part of the data matrix is observed, while in logistic PCA, all matrix entries are observed. Second, logistic PCA and 1-bit matrix completion have somewhat different goals. The main purpose of the former is dimension reduction, while the latter is prediction of the most likely values at the unobserved entries.

Maximum likelihood algorithms (Collins, Dasgupta, and Schapire 2001; Schein, Saul, and Ungar 2003; De Leeuw 2006) have been developed for logistic PCA, which share the same objective function (3). These approaches parameterize Θ with an explicit low-rank factorization, that is, $\Theta = \mathbf{W}\mathbf{H}^T$. In particular, De Leeuw (2006) developed an MM algorithm for logistic PCA, which employed the same majorization of the objective (3) as in our work. At each MM iteration, the corresponding rank-constrained least squares problem was solved using a weighted singular value decomposition (SVD). In principle, his approach may be adapted to the 1-bit matrix completion problem. However, as the simulations in the supplementary materials demonstrate, when the number of observed entries is small, our method is substantially faster.

Binary matrix factorization. The goal of binary matrix factorization (BMF) is to recover a latent low-rank matrix of probabilities from a binary matrix that is typically fully observed (Kabán and Bingham 2008; Bingham, Kabán, and Fortelius 2009; Lumberras, Filstroff, and Févotte 2020; Magron and Févotte 2022). In the BMF problem, the entries are independent draws from a Bernoulli distribution, that is, $y_{ij} \sim \text{Bernoulli}([\mathbf{W}\mathbf{H}^T]_{ij})$ where $\mathbf{W} \in \mathbb{R}^{m \times r}$ and $\mathbf{H} \in \mathbb{R}^{n \times r}$. That is, BMF assumes an identity link function $\Phi(\theta) = \theta$. This is in contrast to a distribution function Φ for the link in 1-bit matrix completion. In BMF, to guarantee that the entries of $\mathbf{W}\mathbf{H}^T$ are probabilities, additional constraints are imposed on the factor matrices, for example, the rows of \mathbf{W} are assumed to be stochastic and the entries of \mathbf{H} binary.

One of the main uses of BMF is dimensionality reduction for exploratory analysis. The identity link is employed because it produces factor matrices with better interpretability. The additional constraints accompanying the identity link, however, lead to a more complicated computational task. Despite having a somewhat different goal, BMF methods could be used to solve the 1-bit matrix completion problem. Simulations in the supplementary materials show that MMGN is typically more accurate and faster than the state-of-the-art BMF method NBMF-MM (Magron and Févotte 2022).

Link prediction in network analysis. In an unweighted undirected network, links between nodes are represented by a binary adjacency matrix. This binary adjacency matrix may be partially observed. A fundamental task in network analysis is to predict whether a link is present at a corresponding unobserved entry in the adjacency matrix (Liben-Nowell and Kleinberg 2003). A

variety of link prediction methods have been developed, among which the generative model-based methods (Miller, Jordan, and Griffiths 2009; Zhou 2015a, 2015b) share similarities with 1-bit matrix completion. To model the binary link y_{ij} between nodes i and j in a network with n nodes, they assume $y_{ij} \sim \text{Bernoulli}(\Phi([\mathbf{Z}\mathbf{W}\mathbf{Z}^T]_{ij}))$ where $\mathbf{Z} \in \mathbb{R}^{n \times k}$, $\mathbf{W} \in \mathbb{R}^{k \times k}$, and Φ is a link function. Assumptions on \mathbf{Z} , \mathbf{W} , and Φ vary across different models. Miller, Jordan, and Griffiths (2009) assumed \mathbf{Z} to be a binary latent feature matrix. Zhou (2015a, 2015b) assumed a Bernoulli-Poisson link $\Phi(\theta) = 1 - e^{-\theta}$, which implies that if the binary observation $y_{ij} = 0$, then the latent mean parameter $[\mathbf{X}\mathbf{W}\mathbf{Z}^T]_{ij} = 0$. Thus, a key difference between link prediction and 1-bit matrix completion is that for some observed matrix entries, there is no uncertainty in the underlying latent mean parameter matrix.

4. Numerical Studies

We compared MMGN¹ to two 1-bit matrix completion methods, TraceNorm (Davenport et al. 2014) and MaxNorm (Cai and Zhou 2013), on simulated data as well as a real data example. We used the computationally more efficient version of TraceNorm that omits the infinity norm constraint $\|\Theta\|_\infty \leq \alpha$. Although Davenport et al. (2014) required the constraint to establish their error bound, they observed that omitting it did not significantly affect the estimation performance and simplified the optimization problem. The methods in Bhaskar and Javanmard (2015) and Ni and Gu (2016) have no publicly available code and consequently were not included in our comparisons. We used MATLAB implementations of TraceNorm and MaxNorm provided by their authors.

In addition, the original problem (2) can be solved by generic optimization approaches. Thus, we also compared MMGN to three general schemes: gradient descent (GD) with backtracking line search as in MMGN, limited-memory BFGS (LBFGS), and Manopt (Boumal et al. 2014) for optimization on the Riemannian manifold of fixed-rank matrices. Details on the implementations appear in the supplementary materials. We also applied GD and LBFGS to ridge regularized version of 1-bit matrix completion and reported results on these variations in the supplementary materials.

4.1. Simulation Settings

Data generation. In our simulations, we considered two types of underlying low-rank matrices: non-spiky and spiky. We generated matrices of each type as follows.

- *Non-spiky.* To generate a non-spiky matrix $\Theta^* \in \mathbb{R}^{m \times n}$ of rank r^* , we constructed $\mathbf{U}^* \in \mathbb{R}^{m \times r^*}$, $\mathbf{V}^* \in \mathbb{R}^{n \times r^*}$ with independent and identically distributed (i.i.d.) entries from a uniform distribution on $[-0.5, 0.5]$. We set $\Theta^* = \mathbf{U}^* \mathbf{V}^{*\top}$ and scaled it so that $\|\Theta^*\|_\infty = 1$. This data generation procedure was the same as in Davenport et al. (2014) and Cai and Zhou (2013).

- *Spiky.* Here, we constructed $\mathbf{U}^* \in \mathbb{R}^{m \times r^*}$, $\mathbf{V}^* \in \mathbb{R}^{n \times r^*}$ with i.i.d. entries from the t -distribution with ν degrees of freedom. We set $\Theta^* = \mathbf{U}^* \mathbf{V}^{*\top}$, without rescaling. A smaller value of ν yielded a heavier-tailed distribution, resulting in a spikier matrix.

For each simulation, we first generated an underlying low-rank matrix Θ^* . Then we randomly selected the set of indices Ω from a uniform distribution with a user-defined observation fraction ρ , namely $\rho = \frac{|\Omega|}{mn}$. Finally, we produced the binary matrix \mathbf{Y} according to the model (1).

Rank estimation. In practice, we need to estimate the rank from the data. In all simulations, we used the following validation approach to estimate the rank. Given a candidate set of K ranks $\{r_1, \dots, r_K\}$, we randomly split the observations into a training set (80%) and a validation set (20%). For each candidate rank r_k , we computed an estimate $\hat{\Theta}_k$ using the training set and then calculated the likelihood of $\hat{\Theta}_k$ on the validation set. We selected the candidate rank with the highest likelihood among all K candidates as our estimated rank r . Even when given a rank r , TraceNorm and MaxNorm could output a matrix with a higher rank. Consequently, we computed a truncated rank- r SVD on their outputs to return a rank- r matrix.

Performance evaluation. We evaluated the methods with the following three metrics.

1. The relative error $\|\hat{\Theta} - \Theta^*\|_F^2 / \|\Theta^*\|_F^2$.
2. The Hellinger distance between the distribution $\Phi(\hat{\Theta})$ and the true distribution $\Phi(\Theta^*)$. Following Davenport et al. (2014) and Cai and Zhou (2013), we define the Hellinger distance between two matrices $\mathbf{P}, \mathbf{Q} \in [0, 1]^{m \times n}$ as

$$d_H^2(\mathbf{P}, \mathbf{Q}) = \frac{1}{mn} \sum_{ij} d_H^2(\mathbf{P}_{ij}, \mathbf{Q}_{ij}),$$

where $d_H^2(p, q) = (\sqrt{p} - \sqrt{q})^2 + (\sqrt{1-p} - \sqrt{1-q})^2$ for $p, q \in [0, 1]$. The Hellinger distance $d_H^2(p, q)$ is a standard metric between two probability distributions. It is nonnegative and equals zero when the two distributions are identical.

It provides a different perspective on estimation accuracy compared to relative error as we will see later at the end of Section 4.3.

3. The runtime in seconds.

Some comparison methods in our simulation studies had fairly long run times. Consequently, we ran only 20 replicates for each simulation setting. Additionally, performance measures for some of the comparison methods exhibited large variability. Consequently, to better visualize trends, the median as well as the 25th and 75th quantiles of metrics are shown using boxplots.

4.2. Non-spiky Matrices

We first compared the performance of different methods when the underlying matrix is non-spiky. We considered the probit noise model. We varied problem parameters, including the noise level σ , the observation fraction ρ , the matrix dimension n , and the true rank r .

¹We implemented MMGN in both Matlab and R. All codes and simulation scripts are available at <https://github.com/Xiaoqian-Liu/MMGN>.

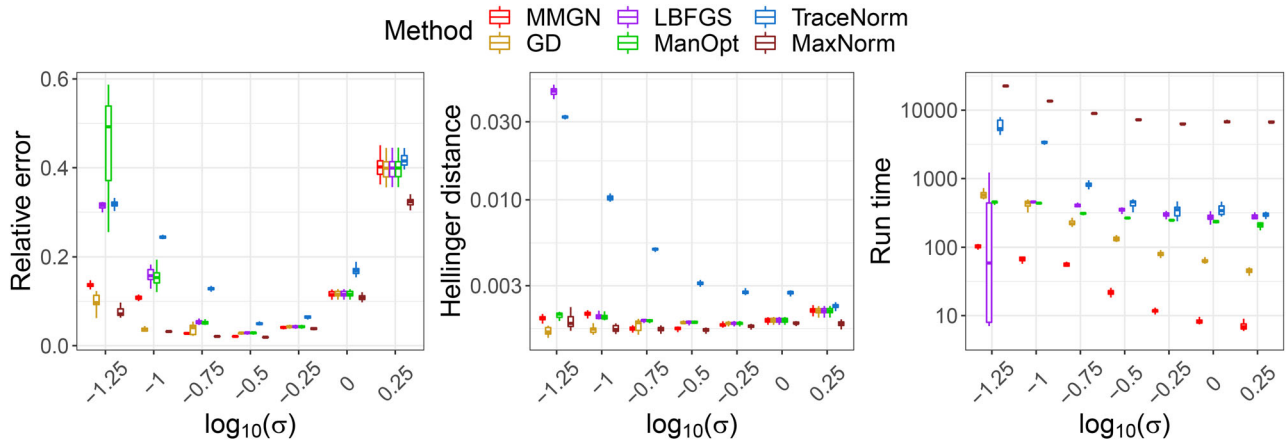


Figure 1. Probit model: Relative error, Hellinger distance, and runtime versus noise level σ for a non-spiky underlying matrix of size $m \times n = 1000 \times 1000$ and rank $r^* = 1$ at observation fraction $\rho = 0.3$.

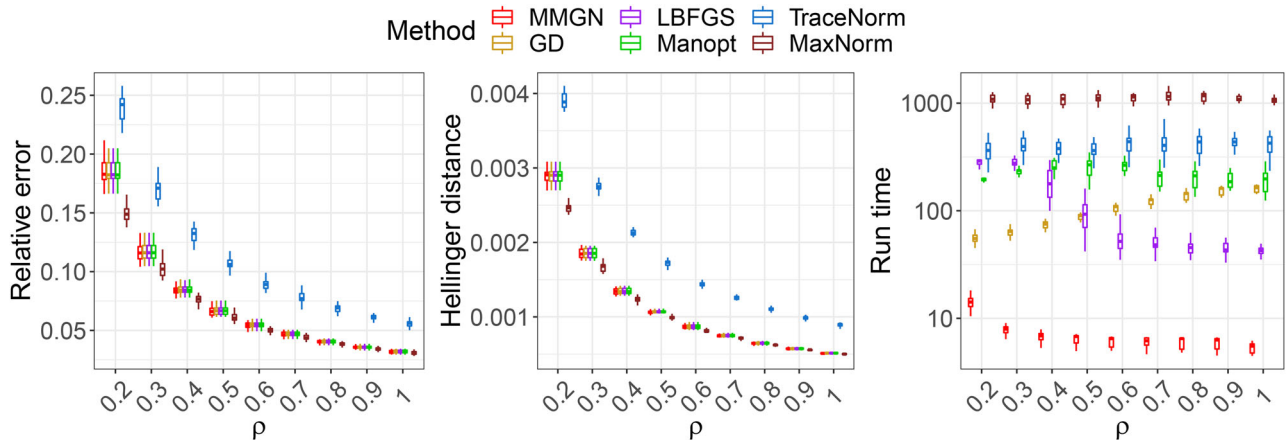


Figure 2. Probit model ($\sigma = 1$): Relative error, Hellinger distance, and runtime versus observation fraction ρ for a non-spiky underlying matrix of size $m \times n = 1000 \times 1000$ and rank $r^* = 1$.

4.2.1. Varying the Noise Level

As discussed in Davenport et al. (2014), the noise level σ is a crucial parameter in 1-bit matrix completion. As σ tends to zero, the problem becomes ill-posed, as two different positive values for θ_{ij}^* yield the same value $y_{ij} = 1$. The first simulation explores the sensitivity of different methods to the noise level σ . We varied σ from $10^{-1.25}$ to $10^{0.25}$ over equally spaced values on a log-linear grid. We set the matrix dimension $m = n = 1000$ and rank $r^* = 1$. The generated underlying matrices are non-spiky with an average spikiness ratio of 3.02 and a standard deviation of 0.07. For each method, we estimated its rank r from $\{1, 2, 3, 4, 5\}$ by the validation procedure described in Section 4.1.

Figure 1 shows the boxplots of relative errors, Hellinger distances, and runtimes (in seconds) of all methods with an observation fraction $\rho = 0.3$. The left panel shows that when the noise level was either too low or too high, all methods produced poor relative errors. This corroborates the results in Davenport et al. (2014) and Cai and Zhou (2013). `MaxNorm` performed slightly better than the other methods when the noise level was extremely low or high. `TraceNorm` suffered from the highest relative error over a wide range of noise levels. `MMGN` and the three generic schemes for solving problem (2) achieved comparable errors when the noise level $\sigma \geq 10^{-0.75}$. The

middle panel shows that the Hellinger distance of `TraceNorm` was significantly larger than those of the other methods, except for $\sigma = 10^{-1.25}$ where the Hellinger distance of `LBFSGS` was even larger. As seen in the left and middle panels of Figure 1, the relative errors and Hellinger distances disagreed in their assessments of the methods. The two metrics capture different errors. We will discuss these differences later. The right panel shows that runtimes of all methods decreased as the noise level grew, and `MMGN` ran the fastest among all compared methods. `MaxNorm` was more than 20 times slower than all the other methods over a wide range of noise levels. In light of this, in all subsequent simulations we supplied `MaxNorm` with the true rank r^* instead of estimating it using the validation procedure.

4.2.2. Varying the Observation Fraction

In the second simulation, we investigated the performance of different methods as a function of the observation fraction $\rho \in \{0.2, 0.3, \dots, 1\}$. The underlying matrices were the same as in the simulations in Section 4.2.1. We considered the probit model with a noise level $\sigma = 1$. We estimated the rank r from $\{1, 2, 3, 4, 5\}$ by the validation procedure described in Section 4.1 for all methods except for `MaxNorm` which was provided with the true rank.

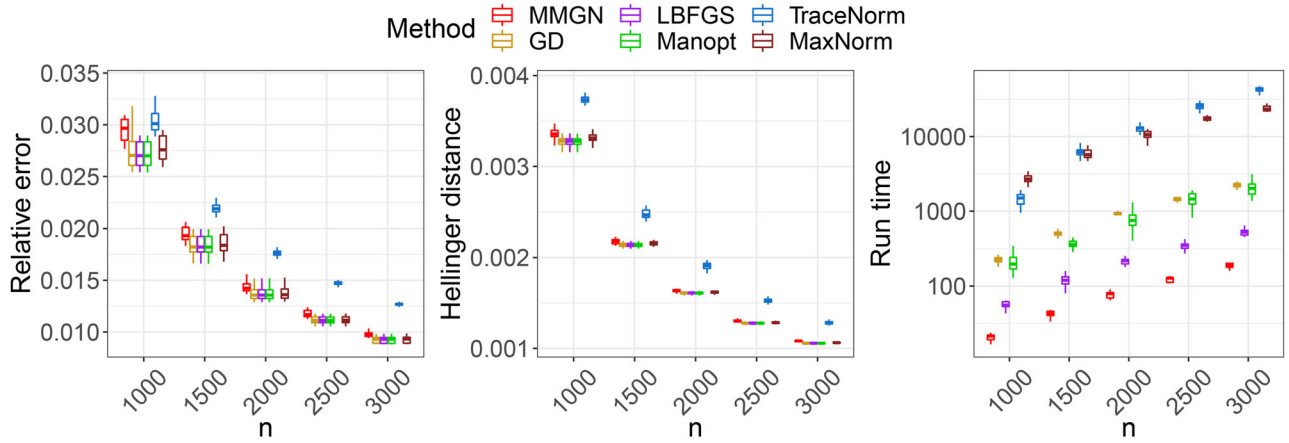


Figure 3. Probit model ($\sigma = 0.18$): Relative error, Hellinger distance, and runtime versus matrix dimension n for a square non-spiky underlying matrix of rank $r^* = 5$ with observation fraction $\rho = 0.8$.




Figure 4. Probit model ($\sigma = 0.18$): Relative error, Hellinger distance, and runtime versus true rank r^* for a non-spiky underlying matrix of size $m \times n = 1000 \times 1000$ and observation fraction $\rho = 0.8$.

The results of this simulation are shown in Figure 2. As expected, the relative errors as well as the Hellinger distances of all methods decreased as the observation fraction ρ increased. Overall, MMGN behaved comparably with MaxNorm and consistently outperformed TraceNorm. Compared to the three generic schemes, GD, LBFGRS, and Manopt, MMGN achieved similar estimation performance while running more than five times faster as seen in the right panel of Figure 2. Moreover, MMGN required substantially less time to achieve comparable estimation accuracy with MaxNorm even when MaxNorm enjoyed the advantage of employing the true rank r^* .

4.2.3. Varying the Matrix Dimension

Next, we compared the performance of different methods under different matrix dimensions. We considered underlying square matrices of dimension $n \in \{1000, 1500, 2000, 2500, 3000\}$, all with rank $r^* = 5$. The average spikiness ratios of the generated matrices were $\{4.61, 4.81, 5.04, 5.09, 5.19, 5.22\}$ for the corresponding values of n . We considered the probit model with $\sigma = 0.18$ and fixed the observation fraction $\rho = 0.8$. For all methods other than MaxNorm which was given the true rank r^* , we employed the validation approach to estimate the rank r from a candidate set $\{3, 4, 5, 6, 7\}$. The results are shown in Figure 3.

Overall, all methods exhibited comparable accuracies that improved as the matrix dimension n grew, with TraceNorm's accuracy being slightly worse than the others. The right panel of Figure 3 shows MMGN's computational advantage over all the other methods. MMGN was up to hundreds of times faster than TraceNorm and MaxNorm. In addition, it typically ran 2–10 times faster than the three generic schemes for solving the optimization problem (2).

4.2.4. Varying the Matrix Rank

In the last simulation under the non-spiky case, we examined how the performance of each method scaled with the rank of the underlying matrix. We fixed the size of Θ^* at $m = n = 1000$ and generated matrices of rank $r^* \in \{3, 5, 8, 10\}$. The average spikiness ratios of the generated underlying matrices for the corresponding values of r^* were $\{4.54, 4.81, 5.03, 4.97\}$. We considered the probit model with a noise level $\sigma = 0.18$ and set the observation fraction $\rho = 0.8$. For all methods except for MaxNorm, we estimated the rank r from $\{1, 2, \dots, 12\}$.

Figure 4 displays the performance of different methods under varying rank r^* . As expected, as r^* increased, the estimation error of each method increased since the number of parameters to be estimated increased. All methods behaved comparably in estimation accuracy, with the Hellinger distance of Tra-

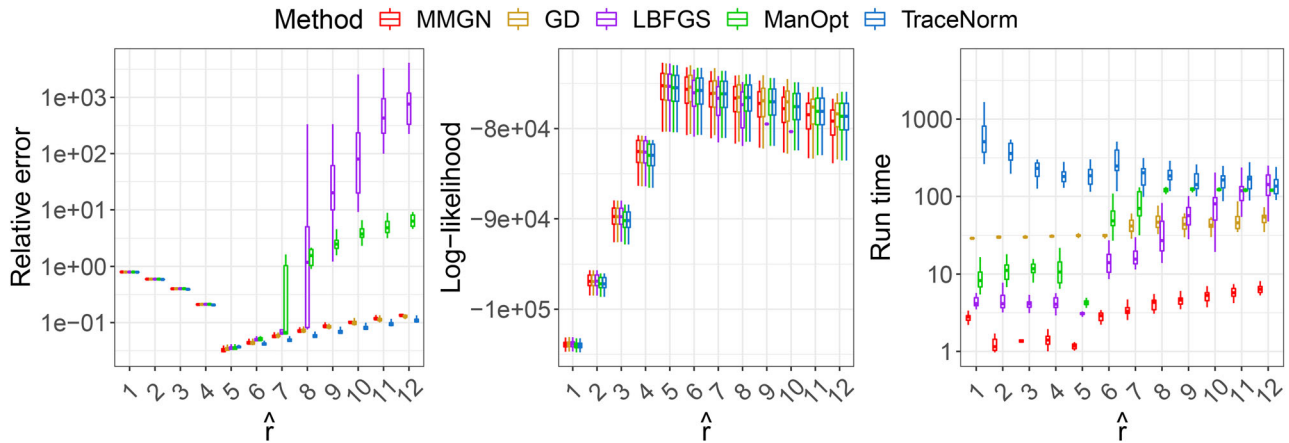


Figure 5. Probit model ($\sigma = 0.18$): Relative error (using the training set), log-likelihood (on the testing set), and runtime versus estimated rank \hat{r} for a non-spiky underlying matrix of size $m \times n = 1000 \times 1000$, true rank $r^* = 5$, and observation fraction $\rho = 0.8$. LBFSG produced log-likelihood of negative infinity at $\hat{r} \in \{11, 12\}$, so the corresponding boxplots were missing.

ceNorm slightly worse than the other methods across different ranks. MMGN enjoyed the fastest runtime among all the compared methods. Recall that we provided the true rank r^* for MaxNorm, thus, its runtime would be even longer if the rank were selected by validation.

To evaluate the rank estimation procedure and illustrate how different methods perform given different estimated ranks, we reran the above simulation with the true rank $r^* = 5$ and reported the relative error using the training set, the log-likelihood on the testing set, and the runtime of each method when given an estimated rank \hat{r} from the set $\{1, 2, \dots, 12\}$. Note that MaxNorm was not included since we ran it without the rank estimation procedure. As shown in Figure 5, all methods achieved the highest log-likelihood as well as the lowest estimation error at the true rank $r^* = 5$. In other words, they selected the correct rank by the procedure. As shown in the right panel of Figure 5, MMGN, GD, LBFSG, and Manopt required longer runtimes when given a rank higher than the true rank. This explains why the runtimes of the four methods in Figure 4 decreased as the true rank r^* increased. With a fixed set of $\hat{r} \in \{1, 2, \dots, 12\}$, as r^* increased from 3 to 10, the number of values of \hat{r} that were higher than r^* decreased, which led to shorter overall runtime.

4.3. Spiky Matrices

We now consider 1-bit matrix completion with spiky underlying matrices, which has not been investigated in the literature. We generated spiky matrices as described in Section 4.1. Specifically, we considered a square matrix $\Theta^* \in \mathbb{R}^{m \times n}$ with $m = n = 1000$ and rank $r^* = 1$. We set the degrees of freedom ν to 10, 5, and 4 to generate underlying matrices at low, intermediate, and high spikiness levels. The corresponding average spikiness ratios (with standard deviation in the parenthesis) over 20 replicates are 17.57 (2.02), 33.16 (4.94), and 47.19 (5.65), respectively. We considered the probit model with a noise level $\sigma = 2$. At each spikiness level, we varied the observation fraction ρ over $\{0.2, 0.3, \dots, 1\}$. We estimated the rank r from $\{1, 2, 3, 4, 5\}$ using the validation procedure for MMGN, TraceNorm, GD, LBFSG and Manopt but provided the true rank for MaxNorm.

Figure 6 shows results under the three spikiness levels. The top panel shows that at the low spikiness level, MMGN and the three generic schemes consistently achieved higher accuracy, in recovering both the underlying matrix Θ^* and the distribution matrix $\Phi(\Theta^*)$, than TraceNorm and MaxNorm over a wide range of values of ρ . MaxNorm achieved smaller errors in comparison to TraceNorm when $\rho < 0.8$ but larger errors at higher observation fractions when $\rho \geq 0.8$. However, MaxNorm recovered the underlying distribution more accurately than TraceNorm, as shown in the top middle panel of Figure 6.

As the spikiness level grew, all methods suffered from lower estimation accuracy, but MMGN and the three generic schemes performed better than TraceNorm and MaxNorm in estimating Θ^* . MaxNorm, MMGN, and the three generic schemes achieved comparable Hellinger distances and significantly outperformed TraceNorm at intermediate and high spikiness levels. Among all compared methods, TraceNorm was the most sensitive to the spikiness level.

An interesting phenomenon is that for all methods except TraceNorm, their errors in estimating Θ^* became relatively insensitive to ρ as the spikiness level increased. The relative errors were dominated by errors in the large-magnitude entries, and errors in these entries were mildly affected by ρ . These methods produced accurate estimates for the majority of low-magnitude entries, even when ρ was small. Therefore, increasing ρ did not lead to significant improvement in their estimation errors. In contrast, TraceNorm required a large value of ρ to produce good estimates for low-magnitude entries. As a result, we saw a more noticeable decrease in its estimation error as ρ grew. We presented a detailed example to elaborate on this phenomenon in the supplementary materials.

The right panel of each row in Figure 6 shows the computational advantage of MMGN when dealing with spiky matrices at different spikiness levels. It ran dozens of times faster than TraceNorm and MaxNorm. Compared to the three generic schemes for the optimization problem (2), MMGN ran typically 2–10 times faster. LBFSG exhibited runtimes comparable with MMGN but only when the observation fraction ρ was close to one.

The simulations so far investigated metrics of overall accuracy defined as the average of the squared errors over all entries

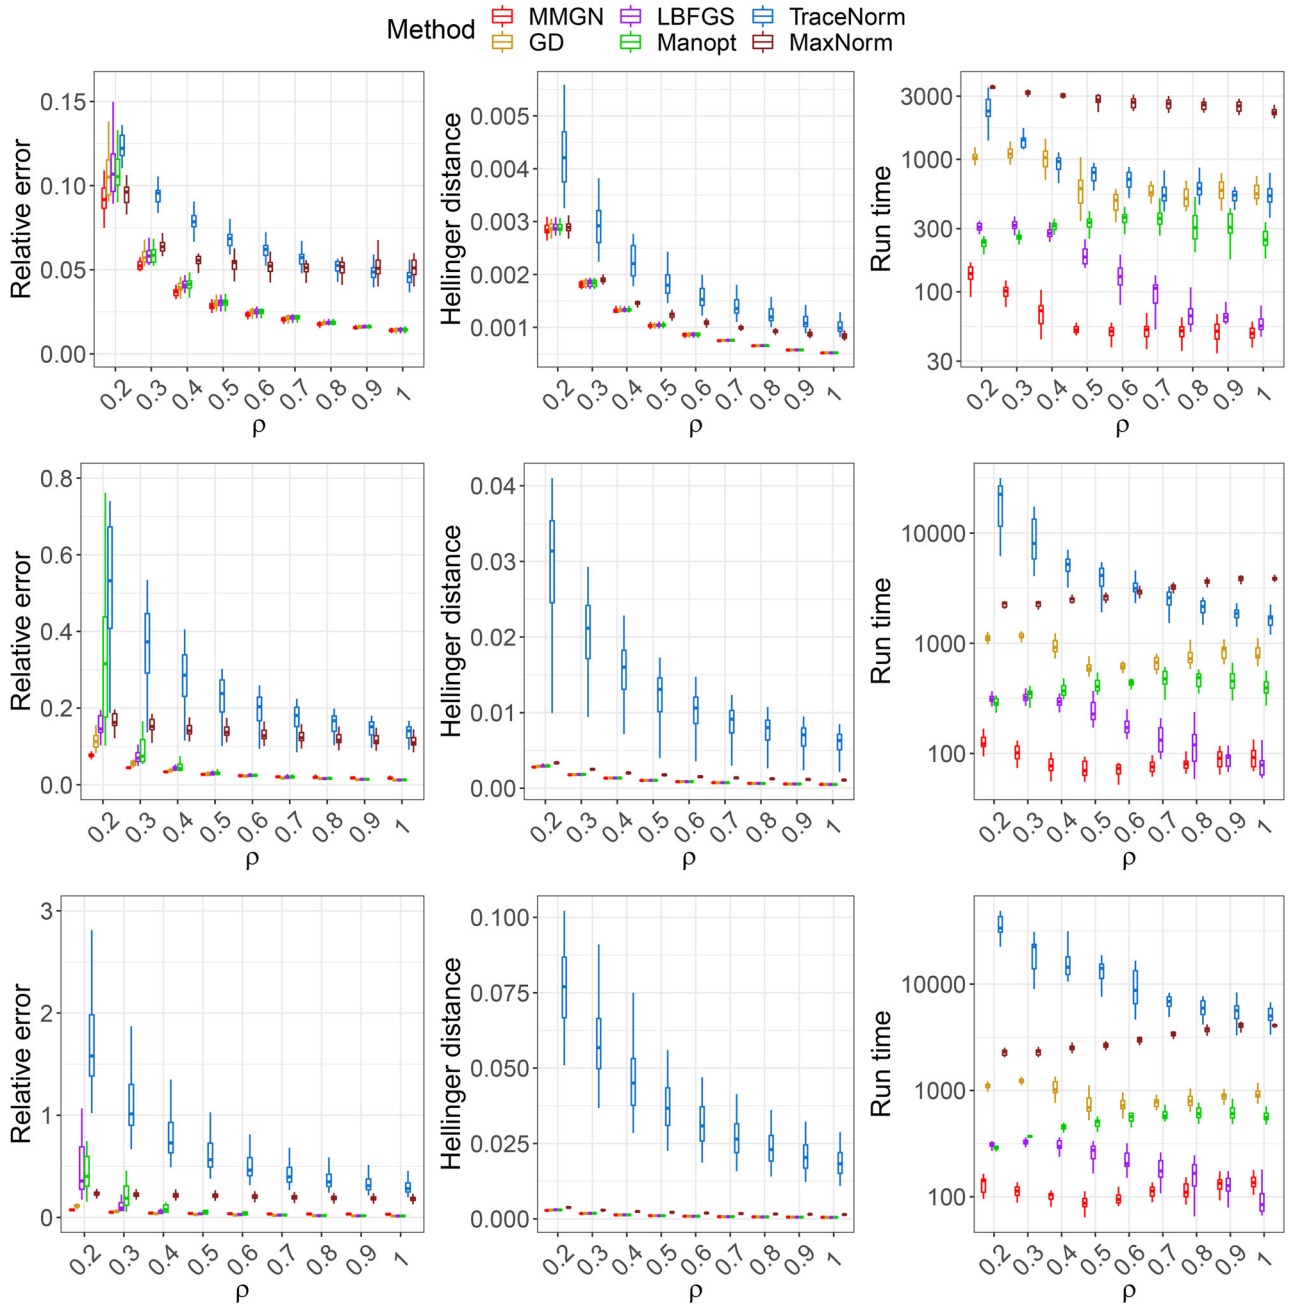


Figure 6. Probit model ($\sigma = 2$): Relative error, Hellinger distance, and runtime versus observation fraction ρ at low (top), intermediate (middle), and high (bottom) spikiness levels for an underlying matrix of size $m \times n = 1000 \times 1000$ and rank $r^* = 1$. Average spikiness ratios are 17.57 (top), 33.16 (middle), and 47.19 (bottom).

of the matrix. However, a more careful inspection reveals that the methods perform differently depending on the magnitude of entries of the underlying matrix Θ^* . To better understand this more fine-grained behavior, we studied the errors conditioned on the underlying values θ_{ij}^* , for a matrix Θ^* with a low spikiness ratio.

Specifically, we considered an underlying matrix Θ^* of size $m \times n = 1000 \times 1000$, rank $r^* = 1$, and spikiness ratio $s = 19.55$. The observation fraction $\rho = 0.8$. The resulting overall relative errors for MMGN, TraceNorm, and MaxNorm were 1.84×10^{-2} , 4.93×10^{-2} , and 5.60×10^{-2} . The corresponding Hellinger distances were 6.30×10^{-4} , 1.10×10^{-3} , and 9.55×10^{-4} , respectively. We divided entries of Θ^* into Value-Ranges that θ_{ij}^* took on. These Value-Ranges have corresponding

Probability-Ranges. For instance, Value-Range $(-\infty, -2.5]$ corresponds to Probability-Range $(0, 0.1]$ since $\lim_{\theta \rightarrow \infty} \Phi(\theta) = 0$ and $\Phi(-2.5) \approx 0.1$ under the considered probit model with $\sigma = 2$.

Figure 7 shows the performance of each method over the different Value- and Probability-Ranges. TraceNorm and MaxNorm yielded accurate estimates for small-magnitude entries of Θ^* (absolute values less than or equal to 2.5) but poorer estimates for large-magnitude entries. MaxNorm suffered from much larger errors in estimating large-magnitude entries, leading to an overall relative error larger than TraceNorm. However, MaxNorm accurately estimated small-magnitude entries leading to accurate recoveries of the corresponding $\Phi(\theta_{ij}^*)$ and a smaller Hellinger distance. This

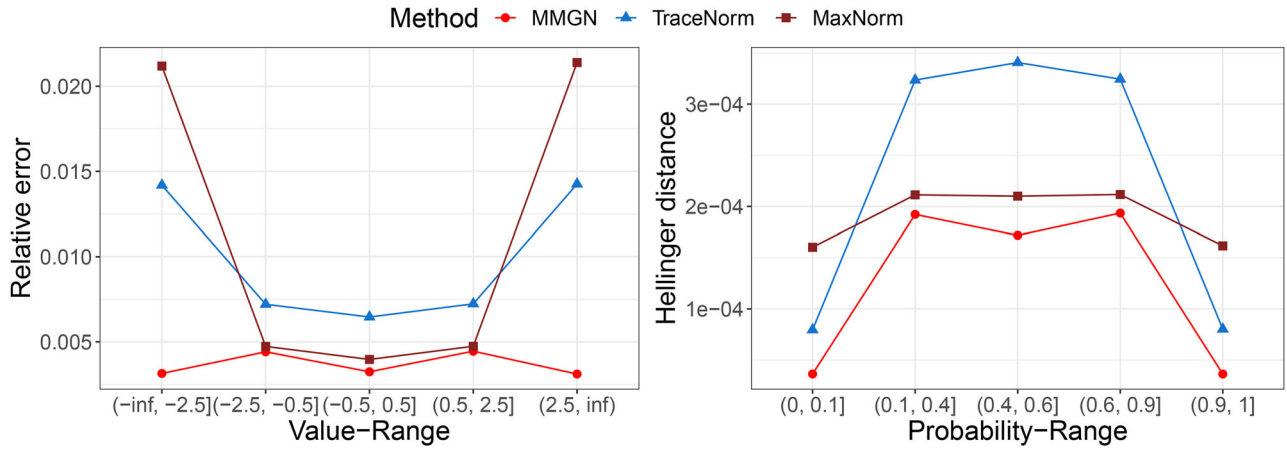


Figure 7. Estimation performance of different methods versus underlying matrix entry values θ_{ij} and probabilities $\Phi(\theta_{ij})$ for a spiky underlying matrix of size $m \times n = 1000 \times 1000$, rank $r^* = 1$, and spikiness ratio $s = 19.55$.

Table 1. Performance of different methods on MovieLens (1M) data.

Method	MMGN	GD	LBFGS	Manopt	TraceNorm	MaxNorm
Accuracy (%)	74.2	74.4	74.0	74.1	75.0	74.0
Time (seconds)	6.7e+2	3.36e+3	2.25e+04	4.08e+03	1.20e+4	6.41e+4

discrepancy is caused by the nonlinear transformation from θ_{ij}^* to $\Phi(\theta_{ij}^*)$. For example, $\Phi(0.5) - \Phi(0.4) \approx 0.019$, while $\Phi(3.6) - \Phi(3.5) \approx 0.004$. In other words, the recovered distribution is more sensitive to the accuracy of the recovered $\hat{\theta}_{ij}$ in small-magnitude groups. In comparison, MMGN achieved the highest accuracy across different Value-Ranges, leading to the best overall performance. Figure 7 provides a more detailed picture of how MMGN outperformed TraceNorm and MaxNorm. Moreover, this explains why the trends in estimating Θ^* do not mirror the trends in estimating $\Phi(\Theta^*)$ seen in Figures 1 and 6.

4.4. Real Data Example

We applied MMGN to the MovieLens (1M) dataset,² which is larger but otherwise similar to the MovieLens (100K) real data example in Davenport et al. (2014). It contains 1,000,209 movie ratings from 6040 users on 3952 movies, with each rating on a scale from 1 to 5. The resulting observation fraction is 0.04. Following Davenport et al. (2014), we converted all ratings to binary observations by comparing each rating to the average rating over all movies. Ratings above the average were encoded as +1 and -1 otherwise. As in Davenport et al. (2014), we considered the logistic model with the noise level $\sigma = 1$ and used 95% of the ratings as the training set to estimate Θ^* . The performance was evaluated on the remaining 5% of ratings by checking whether or not the estimate of Θ^* correctly predicts the sign of the ratings (above or below the average rating). We compared MMGN to TraceNorm, MaxNorm, GD, LBFGS, and Manopt. We followed Davenport et al. (2014) and treated $\alpha\beta$ as a single parameter in TraceNorm. We assigned $\alpha\beta$ 10 values equally spaced on a logarithmic scale between $10^{-0.5}$ to 10. We ran MMGN and the three generic schemes with a candidate rank

r ranging from 1 to 10. For MaxNorm, β was set to \sqrt{r} , where again r ranged from 1 to 10, and α was chosen such that $\alpha\beta$ took the same value as $\alpha\beta$ for TraceNorm. As in Davenport et al. (2014), we selected parameter values for each method that led to the best prediction performance on the separate validation set.

Table 1 summarizes the median prediction accuracies and runtimes of the compared methods over 20 replicates. All methods achieved comparable prediction accuracies, but MMGN was more than an order of magnitude faster than the two existing 1-bit matrix completion methods and at least 5-fold faster than the three generic schemes. This example highlights the computational advantage of MMGN in dealing with large sparse (very few observations) datasets from real-world problems.

4.5. Additional Comparisons

As mentioned above, 1-bit matrix completion can be viewed as a logistic PCA problem but with only partially observed entries. In particular, one could solve the 1-bit objective by a logistic-PCA scheme that is able to handle missing data. De Leeuw (2006) proposed such a scheme, denoted logisticPCA. Thus, we compared MMGN to logisticPCA. Interestingly, this scheme is also based on MM. The difference between the two algorithms lies in how they solve the rank-constrained least squares problem. The per-iteration computation in logisticPCA is dominated by a rank- r truncated SVD. The implementation we used for logisticPCA is based on ARPACK which requires r rounds of the Lanczos method for a total cost of $O(mnr)$ flops; each round of the Lanczos method costs $O(mn)$ flops, due to matrix-vector multiplies involving an m -by- n matrix (Lehoucq, Sorensen, and Yang 1998). In contrast, MMGN adopts a single Gauss-Newton step to approximate the rank-constrained least squares problem by a sparse least square problem whose computational cost is $\mathcal{O}(|\Omega|r)$. A detailed breakdown of this cost is provided in the supplementary materials. In simulations, which

²The dataset is available at <http://www.grouplens.org/node/73>.

can be found in the supplementary materials, `logisticPCA` typically achieved accuracies comparable with `MMGN` but ran 2–10 times slower than `MMGN`, especially in the more challenging cases where the observation fraction is small.

We also compared `MMGN` to the nonnegative binary matrix factorization (`NBMF-MM`) method proposed by Magron and Févotte (2022). `NBMF-MM` is the state-of-the-art method for BMF, displaying computational advantages over competing BMF methods (Kabán and Bingham 2008; Bingham, Kabán, and Fortelius 2009; Lumbreras, Filstroff, and Févotte 2020), as discussed in Magron and Févotte (2022). Given a target rank of r , the per-iteration computational cost of `NBMF-MM` is $\mathcal{O}(mnr)$ flops. In simulations, which can be found in the supplementary materials, `MMGN` was more accurate than `NBMF-MM`, while running typically dozens of times faster.

We also compared `MMGN` with a 1-bit tensor completion (`1BitTC`) method (Wang and Li 2020) viewing our matrix as a three-way tensor with the third mode having dimension one. While several methods were proposed for binary tensor completion (Li et al. 2018; Aidini, Tsagkatakis, and Tsakalides 2018; Ghadermarzy, Plan, and Yilmaz 2018; Wang and Li 2020), we focused on `1BitTC` since, to the best of our knowledge, it has the best error rates, as well as publicly available code. Since `1BitTC` works with tensors, it is expected to be much slower than methods that work directly with matrices. As shown in the supplementary materials, in addition to being slower, `1BitTC` produced higher errors under many scenarios compared with `MMGN`.

5. Summary and Discussion

In this article, we proposed `MMGN`, a novel fast and accurate algorithm for solving the 1-bit optimization problem. Our algorithm employs the MM principle to convert the original challenging problem (2) into a sequence of standard low-rank matrix completion problems. For each MM update, we apply a factorization strategy to incorporate the low-rank structure, resulting in a nonlinear least squares problem. We then apply a modified Gauss-Newton scheme to compute an inexact MM update by solving a linear least squares problem. Hence, in comparison to previous works, our method is relatively simple and can easily scale to large matrices. In comparison to two established 1-bit matrix completion methods, `TraceNorm` (Davenport et al. 2014) and `MaxNorm` (Cai and Zhou 2013), and three generic schemes for solving the original optimization problem (2), `MMGN` typically achieved comparable and sometimes lower errors while often being significantly faster.

Our work raises several interesting directions for further research. On the practical side, it is of interest to extend `MMGN` to other quantization models, such as movie ratings between 1 and 5. On the theoretical side, an open problem is to prove that under suitable assumptions, and possibly starting from a sufficiently accurate initial guess, `MMGN` indeed converges to the maximum likelihood solution.

Supplementary Materials

Title: Supplementary Materials for “A Majorization-Minimization Gauss-Newton Method for 1-Bit Matrix Completion.” (.pdf file)

Acknowledgments

B.N. is incumbent of the William Petschek Professorial Chair of Mathematics. We thank Wenxin Zhou for sharing us the code of `MaxNorm`. We thank Mark Davenport for helpful discussions. We are grateful to the editor, the associate editor, and two anonymous reviewers for their insightful and constructive comments for improving the quality of the article.

Disclosure Statement

The authors report there are no competing interests to declare.

Funding

The research of B.N. was funded by the National Institutes of Health (NIH) grant R01GM135928 and by the Israel Science Foundation (ISF) grant 2362/22. The research of E.C. was supported by the NIH grant R01GM135928.

ORCID

Eric C. Chi  <http://orcid.org/0000-0003-4647-0895>

References

- Aidini, A., Tsagkatakis, G., and Tsakalides, P. (2018), “1-bit Tensor Completion,” *Electronic Imaging*, 30, 261-1–261-6. [12]
- Bauch, J., Nadler, B., and Zilber, P. (2021), “Rank 2r Iterative Least Squares: Efficient Recovery of Ill-conditioned Low Rank Matrices from Few Entries,” *SIAM Journal on Mathematics of Data Science*, 3, 439–465. [4]
- Bhaskar, S. A., and Javanmard, A. (2015), “1-bit Matrix Completion under Exact Low-Rank Constraint,” in *2015 49th Annual Conference on Information Sciences and Systems (CISS)*, IEEE, pp. 1–6. [2,6]
- Bingham, E., Kabán, A., and Fortelius, M. (2009), “The Aspect Bernoulli Model: Multiple Causes of Presences and Absences,” *Pattern Analysis and Applications*, 12, 55–78. [5,12]
- Boumal, N., Mishra, B., Absil, P.-A., and Sepulchre, R. (2014), “Manopt, a Matlab Toolbox for Optimization on Manifolds,” *Journal of Machine Learning Research*, 15, 1455–1459. [6]
- Cai, T., and Zhou, W.-X. (2013), “A Max-Norm Constrained Minimization Approach to 1-bit Matrix Completion,” *The Journal of Machine Learning Research*, 14, 3619–3647. [2,6,7,12]
- Collins, M., Dasgupta, S., and Schapire, R. E. (2001), “A Generalization of Principal Components Analysis to the Exponential Family,” in *Advances in Neural Information Processing Systems* (Vol. 14). [5]
- Davenport, M. A., Plan, Y., Van Den Berg, E., and Wootters, M. (2014), “1-bit Matrix Completion,” *Information and Inference: A Journal of the IMA*, 3, 189–223. [1,2,6,7,11,12]
- De Leeuw, J. (1994), “Block-Relaxation Algorithms in Statistics,” in *Information Systems and Data Analysis*, eds. H.-H. Bock, W. Lenski, and M. M. Richter, pp. 308–324, Berlin, Heidelberg: Springer. [2]
- (2006), “Principal Component Analysis of Binary Data by Iterated Singular Value Decomposition,” *Computational Statistics & Data Analysis*, 50, 21–39. [5,11]
- Ghadermarzy, N., Plan, Y., and Yilmaz, O. (2018), “Learning Tensors from Partial Binary Measurements,” *IEEE Transactions on Signal Processing*, 67, 29–40. [12]
- Gross, D., Liu, Y.-K., Flammia, S. T., Becker, S., and Eisert, J. (2010), “Quantum State Tomography via Compressed Sensing,” *Physical Review Letters*, 105, 150401. [1]
- Heiser, W. J. (1995), “Convergent Computation by Iterative Majorization,” in *Recent Advances in Descriptive Multivariate Analysis*, ed. W. J. Krzanowski, 157–189, Oxford: Oxford University Press. [2]
- Hestenes, M. R., and Stiefel, E. (1952), “Methods of Conjugate Gradients for Solving Linear Systems,” *Journal of Research of the National Bureau of Standards*, 49, 409–436. [4]
- Hunter, D. R., and Lange, K. (2004), “A Tutorial on MM Algorithms,” *The American Statistician*, 58, 30–37. [2,4]

- Kabán, A., and Bingham, E. (2008), “Factorisation and Denoising of 0–1 Data: A Variational Approach,” *Neurocomputing*, 71, 2291–2308. [5,12]
- Kammerer, W. J., and Nashed, M. Z. (1972), “On the Convergence of the Conjugate Gradient Method for Singular Linear Operator Equations,” *SIAM Journal on Numerical Analysis*, 9, 165–181. [4]
- Koren, Y., Bell, R., and Volinsky, C. (2009), “Matrix Factorization Techniques for Recommender Systems,” *Computer*, 42, 30–37. [1]
- Lange, K. (2016), *MM Optimization Algorithms*, Philadelphia, PA: SIAM. [2]
- Lange, K., Hunter, D. R., and Yang, I. (2000), “Optimization Transfer Using Surrogate Objective Functions,” *Journal of Computational and Graphical Statistics*, 9, 1–20. [2]
- Lehoucq, R. B., Sorensen, D. C., and Yang, C. (1998), *ARPACK Users’ Guide - Solution of Large-Scale Eigenvalue Problems with Implicitly Restarted Arnoldi methods*, Philadelphia, PA: SIAM. [11]
- Li, B., Zhang, X., Li, X., and Lu, H. (2018), “Tensor Completion from One-Bit Observations,” *IEEE Transactions on Image Processing*, 28, 170–180. [12]
- Liben-Nowell, D., and Kleinberg, J. (2003), “The Link Prediction Problem for Social Networks,” in *Proceedings of the Twelfth International Conference on Information and Knowledge Management*, pp. 556–559. [1,5]
- Linial, N., Mendelson, S., Schechtman, G., and Shraibman, A. (2007), “Complexity Measures of Sign Matrices,” *Combinatorica*, 27, 439–463. [2]
- Lumbreras, A., Filstroff, L., and Févotte, C. (2020), “Bayesian Mean-Parameterized Nonnegative Binary Matrix Factorization,” *Data Mining and Knowledge Discovery*, 34, 1898–1935. [5,12]
- Magron, P., and Févotte, C. (2022), “A Majorization-Minimization Algorithm for Nonnegative Binary Matrix Factorization,” *IEEE Signal Processing Letters*, 29, 1526–1530. [5,12]
- Miller, K., Jordan, M., and Griffiths, T. (2009), “Nonparametric Latent Feature Models for Link Prediction,” *Advances in Neural Information Processing Systems* (Vol. 22). [6]
- Negahban, S., and Wainwright, M. J. (2012), “Restricted Strong Convexity and Weighted Matrix Completion: Optimal Bounds with Noise,” *The Journal of Machine Learning Research*, 13, 1665–1697. [1]
- Ni, R., and Gu, Q. (2016), “Optimal Statistical and Computational Rates for One Bit Matrix Completion,” in *Artificial Intelligence and Statistics* (Vol. 51), pp. 426–434, PMLR. [2,6]
- Nocedal, J., and Wright, S. J. (2006), *Numerical Optimization* (2nd ed.), New York, NY: Springer. [4]
- Paige, C. C., and Saunders, M. A. (1982), “LSQR: An Algorithm for Sparse Linear Equations and Sparse Least Squares,” *ACM Transactions on Mathematical Software*, 8, 43–71. [4]
- Schein, A. I., Saul, L. K., and Ungar, L. H. (2003), “A Generalized Linear Model for Principal Component Analysis of Binary Data,” in *International Workshop on Artificial Intelligence and Statistics*, PMLR, pp. 240–247. [5]
- Wang, M., and Li, L. (2020), “Learning from Binary Multiway Data: Probabilistic Tensor Decomposition and its Statistical Optimality,” *Journal of Machine Learning Research*, 21, 6146–6183. [12]
- Xu, J., Chi, E. C., Yang, M., and Lange, K. (2018), “A Majorization–Minimization Algorithm for Split Feasibility Problems,” *Computational Optimization and Applications*, 71, 795–828. [4]
- Zhou, M. (2015a), “Infinite Edge Partition Models for Overlapping Community Detection and Link Prediction,” in *Artificial Intelligence and Statistics*, PMLR, pp. 1135–1143. [6]
- (2015b), “Nonparametric Bayesian Matrix Factorization for Assortative Networks,” in *2015 23rd European Signal Processing Conference (EUSIPCO)*, IEEE, pp. 2776–2780. [6]
- Zilber, P., and Nadler, B. (2022), “GNMR: A Provable One-Line Algorithm for Low Rank Matrix Recovery,” *SIAM Journal on Mathematics of Data Science*, 4, 909–934. [2,3,4]