1-Bit Matrix Completion

Mark A. Davenport, Yaniv Plan, Ewout van den Berg, Mary Wootters*

September 18, 2012

Abstract

In this paper we develop a theory of matrix completion for the extreme case of noisy 1-bit observations. Instead of observing a subset of the real-valued entries of a matrix $\boldsymbol{M}$, we obtain a small number of binary (1-bit) measurements generated according to a probability distribution determined by the real-valued entries of $\boldsymbol{M}$. The central question we ask is whether or not it is possible to obtain an accurate estimate of $\boldsymbol{M}$ from this data. In general this would seem impossible, but we show that the maximum likelihood estimate under a suitable constraint returns an accurate estimate of $\boldsymbol{M}$ when $\|\boldsymbol{M}\|_\infty \le \alpha$ and $\operatorname{rank}(\boldsymbol{M}) \le r$. If the log-likelihood is a concave function (e.g., the logistic or probit observation models), then we can obtain this maximum likelihood estimate by optimizing a convex program. In addition, we also show that if instead of recovering $\boldsymbol{M}$ we simply wish to obtain an estimate of the distribution generating the 1-bit measurements, then we can eliminate the requirement that $\|\boldsymbol{M}\|_\infty \le \alpha$. For both cases, we provide lower bounds showing that these estimates are near-optimal.

1 Introduction

The problem of recovering a matrix from an incomplete sampling of its entries—also known as *matrix completion*—arises in a wide variety of practical situations. In many of these settings, however, the observations are not only incomplete, but also highly *quantized*, often even to a single bit. In this paper we consider a statistical model for such data where instead of observing a real-valued entry as in the original matrix completion problem, we are now only able to see a positive or negative rating. This binary output is generated according to a probability distribution which is parameterized by the corresponding entry of the unknown matrix $\boldsymbol{M}$. The central question we ask in this paper is: "Given observations of this form, can we recover the underlying matrix?"

We will see that $O(rd)$ measurements are sufficient to accurately recover a $d \times d$, rank-r matrix from such data. Before describing this result and others in more detail, we provide a brief review of the matrix completion problem and the closely related problem of 1-bit compressed sensing.

1.1 Matrix completion

Matrix completion arises in a wide variety of practical contexts, including collaborative filtering [17], system identification [32], sensor localization [3, 44, 45], rank aggregation [16], and many more. While many of these applications have a relatively long history, recent advances in the closely related field of compressed sensing [14, 7, 13] have enabled a burst of progress in the last few years, and we now have a strong base of theoretical results concerning matrix completion [19, 10, 11, 24, 25, 36, 29, 9, 41, 42, 27, 15, 26, 28].

*M.A. Davenport is with the School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA. Email: {mdav@gatech.edu}.

Y. Plan is with the Department of Mathematics, University of Michigan, Ann Arbor, MI. Email: {yplan@umich.edu}.

E. van den Berg is with the Department of Statistics, Stanford University, Stanford, CA. Email: {ewout@stanford.edu}.

M. Wootters is with the Department of Mathematics, University of Michigan, Ann Arbor, MI. Email: {wootters@umich.edu}.

This work was partially supported by NSF grants DMS-1004718, DMS-1103909 and CCF-0743372.

A typical result from this literature is that a generic $d \times d$ matrix of rank r can be exactly recovered from $O(r\, d \operatorname{polylog}(d))$ randomly chosen entries. Similar results can be established in the case of noisy observations and approximately low-rank matrices [25, 36, 29, 9, 42, 27, 15, 26, 28].

Although these results are quite impressive, there is an important gap between the statement of the problem as considered in the matrix completion literature and many of the most common applications discussed therein. As an example, consider collaborative filtering and the now-famous "Netflix problem." In this setting, we assume that there is some unknown matrix whose entries each represent a rating for a particular user on a particular movie. Since any user will rate only a small subset of possible movies, we are only able to observe a small fraction of the total entries in the matrix, and our goal is to infer the unseen ratings from the observed ones. If the rating matrix is low-rank, then this would seem to be the exact problem studied in the matrix completion literature. However, there is a subtle difference: the theory developed in this literature generally assumes that observations consist of (possibly noisy) continuous-valued entries of the matrix, whereas in the Netflix problem the observations are "quantized" to the set of integers between 1 and 5. If we believe that it is possible for a user's true rating for a particular movie to be, for example, 4.5, then we must account for the impact of this "quantization noise" on our recovery. Of course, one could potentially treat quantization simply as a form of bounded noise, but this is somewhat unsatisfying because the ratings aren't just quantized — there are also hard limits placed on the minimum and maximum allowable ratings. (Why should we suppose that a movie given a rating of 5 could not have a true underlying rating of 6 or 7 or 10?) The inadequacy of standard matrix completion techniques in dealing with this effect is particularly pronounced when we consider recommender systems where each rating consists of a single bit representing a positive or negative rating (consider for example rating music on Pandora, the relevance of advertisements on Hulu, or posts on Reddit or MathOverflow). In such a case, the assumptions made in the existing theory of matrix completion do not apply, standard algorithms are ill-posed, and a new theory is required.

1.2 1-Bit compressed sensing and sparse logistic regression

As noted above, matrix completion is closely related to the field of compressed sensing, where a theory to deal with single-bit quantization has recently been developed [5, 22, 37, 38, 21, 30]. In compressed sensing, one can recover an s-sparse vector in $\mathbb{R}^d$ from $O(s \log(d/s))$ random linear measurements—several different random measurement structures are compatible with this theory. In 1-bit compressed sensing, only the signs of these measurements are observed, but an s-sparse signal can still be approximately recovered from the same number of measurements [22, 37, 38, 1]. However, the only measurement ensembles which are currently known to give such guarantees are Gaussian or sub-Gaussian [1], and thus of a quite different flavor than the kinds of samples obtained in the matrix completion setting. A similar theory is available for the closely related problem of sparse binomial regression, which considers more classical statistical models [2, 6, 38, 23, 33, 35, 40, 47] and allows non-Gaussian measurements. Our aim here is to develop results for matrix completion of the same flavor as 1-bit compressed sensing and sparse logistic regression.

1.3 Challenges

In this paper, we extend the theory of matrix completion to the case of 1-bit observations. We consider a general observation model but focus mainly on two particular possibilities: the models of logistic and probit regression. We discuss these models in greater detail in Section 2.1, but first we note that several new challenges arise when trying to leverage results in 1-bit compressed sensing and sparse logistic regression to develop a theory for 1-bit matrix completion. First, matrix completion is in some sense a more challenging problem than compressed sensing. Specifically, some additional difficulty arises because the set of low-rank matrices is "coherent" with single entry measurements (see [19]). In particular, the sampling operator does not act as a near-isometry on all matrices of interest, and thus the natural analogue to the restricted isometry property from compressed sensing cannot hold in general—there will always be certain low-rank matrices that we cannot hope to recover without essentially sampling every entry of the matrix. For example, consider a matrix that consists of a single nonzero entry (which we might never observe). The typical way to deal with this possibility is to consider a reduced set of low-rank matrices by placing restrictions on the entry-wise maximum of the matrix or its singular vectors—informally, we require that the matrix is not too "spiky".

We introduce an entirely new dimension of ill-posedness by restricting ourselves to 1-bit observations. To illustrate this, we describe one version of 1-bit matrix completion in more detail (the general problem definition is given in Section 2.1 below). Consider a $d \times d$ matrix $\boldsymbol{M}$ with rank r. Suppose we observe a subset Ω of entries of a matrix $\boldsymbol{Y}$. The entries of $\boldsymbol{Y}$ depend on $\boldsymbol{M}$ in the following way:

$$Y_{i,j} = \begin{cases} +1 & \text{if } M_{i,j} + Z_{i,j} \geq 0 \\ -1 & \text{if } M_{i,j} + Z_{i,j} < 0 \end{cases} \tag{1}$$

where $\boldsymbol{Z}$ is a matrix containing noise. This latent variable model is the direct analogue to the usual 1-bit compressed sensing observation model. In this setting, we view the matrix $\boldsymbol{M}$ as more than just a parameter of the distribution of $\boldsymbol{Y}$; $\boldsymbol{M}$ represents the real underlying quantity of interest that we would like to estimate. Unfortunately, in what would seem to be the most benign setting—when Ω is the set of all entries, $\boldsymbol{Z} = \boldsymbol{0}$, and $\boldsymbol{M}$ has rank 1 and a bounded entry-wise maximum—the problem of recovering $\boldsymbol{M}$ is ill-posed. To see this, let $\boldsymbol{M} = \boldsymbol{u}\boldsymbol{v}^*$ for any vectors $\boldsymbol{u}, \boldsymbol{v} \in \mathbb{R}^d$, and for simplicity assume that there are no zero entries in $\boldsymbol{u}$ or $\boldsymbol{v}$. Now let $\widetilde{\boldsymbol{u}}$ and $\widetilde{\boldsymbol{v}}$ be any vectors with the same sign pattern as $\boldsymbol{u}$ and $\boldsymbol{v}$ respectively. It is apparent that either $\boldsymbol{M}$ or $\widetilde{\boldsymbol{M}} = \widetilde{\boldsymbol{u}}\widetilde{\boldsymbol{v}}^*$ will yield the same observations $\boldsymbol{Y}$, and thus $\boldsymbol{M}$ and $\widetilde{\boldsymbol{M}}$ are indistinguishable. Note that while it is obvious that this 1-bit measurement process will destroy any information we have regarding the scaling of $\boldsymbol{M}$, this ill-posedness remains even if we knew something about the scaling a priori (such as the Frobenius norm of $\boldsymbol{M}$). For any given set of observations, there will always be radically different possible matrices that are all consistent with observed measurements.

After considering this example, the problem might seem hopeless. However, an interesting surprise is that when we add noise to the problem (that is, when $\boldsymbol{Z} \neq \boldsymbol{0}$ is an appropriate stochastic matrix) the picture completely changes—this noise has a "dithering" effect and the problem becomes well-posed. In fact, we will show that in this setting we can sometimes recover $\boldsymbol{M}$ to the same degree of accuracy that is possible when given access to completely unquantized measurements! In particular, under appropriate conditions, $O(rd)$ measurements are sufficient to accurately recover $\boldsymbol{M}$.

1.4 Applications

The problem of 1-bit matrix completion arises in nearly every application that has been proposed for "unquantized" matrix completion. To name a few:

- **Recommender systems:** As mentioned above, collaborative filtering systems often involve discretized recommendations [17]. In many cases, each observation will consist simply of a "thumbs up" or "thumbs down" thus delivering only 1 bit of information (consider for example rating music on Pandora, the relevance of advertisements on Hulu, or posts on Reddit or MathOverflow). Such cases are a natural application for 1-bit matrix completion.

- **Analysis of survey data:** Another potential application for matrix completion is to analyze incomplete survey data. Such data is almost always heavily quantized since people are generally not able to distinguish between more than 7 ± 2 categories [34]. 1-bit matrix completion provides a method for analyzing incomplete (or potentially even complete) survey designs containing simple yes/no or agree/disagree questions.

- **Distance matrix recovery and multidimensional scaling:** Yet another common motivation for matrix completion is to localize nodes in a sensor network from the observation of just a few inter-node distances [3, 44, 45]. This is essentially a special case of multidimensional scaling (MDS) from incomplete data [4]. In general, work in the area assumes real-valued measurements. However, in the sensor network example (as well as many other MDS scenarios), the measurements may be very coarse and might only indicate whether the nodes are within or outside of some communication range. While there is some existing work on MDS using binary data [18] and MDS using incomplete observations with other kinds of non-metric data [46], 1-bit matrix completion promises to provide a principled and unifying approach to such problems.

- **Quantum state tomography:** Low-rank matrix recovery from incomplete observations also has applications to quantum state tomography [20]. In this scenario, mixed quantum states are represented as Hermitian matrices with nuclear norm equal to 1. When the state is nearly pure, the matrix can be well approximated by a low-rank matrix and, in particular, fits the model given in Section 2.2 up to a rescaling. Furthermore, Pauli-operator-based measurements give probabilistic binary outputs. However, these are based on the inner products with the Pauli-matrices, and thus of a slightly different flavor than the measurements considered in this paper. Nevertheless, while we do not address this scenario directly, our theory of 1-bit matrix completion could easily be adapted to quantum state tomography.

1.5 Notation

We now provide a brief summary of some of the key notation used in this paper. We use $[d]$ to denote the set of integers $\{1, \ldots, d\}$. We use capital boldface to denote a matrix (e.g., $\boldsymbol{M}$) and standard text to denote its entries (e.g., $M_{i,j}$). Similarly, we let $\mathbf{0}$ denote the matrix of all-zeros and $\mathbf{1}$ the matrix of all-ones. We let $\|\boldsymbol{M}\|$ denote the operator norm of $\boldsymbol{M}$, $\|\boldsymbol{M}\|_F = \sqrt{\sum_{i,j} M_{i,j}^2}$ denote the Frobenius norm of $\boldsymbol{M}$, $\|\boldsymbol{M}\|_*$ denote the nuclear or Schatten-1 norm of $\boldsymbol{M}$ (the sum of the singular values), and $\|\boldsymbol{M}\|_\infty = \max_{i,j} |M_{i,j}|$ denote the entry-wise infinity-norm of $\boldsymbol{M}$. We will use the Hellinger distance, which, for two scalars $p, q \in [0,1]$, is given by

$$d_H^2(p,q) := (\sqrt{p} - \sqrt{q})^2 + (\sqrt{1-p} - \sqrt{1-q})^2.$$

This gives a standard notion of distance between two binary probability distributions. We also allow the Hellinger distance to act on matrices via the average Hellinger distance over their entries: for matrices $\boldsymbol{P}, \boldsymbol{Q} \in [0,1]^{d_1 \times d_2}$, we define

$$d_H^2(\boldsymbol{P}, \boldsymbol{Q}) = \frac{1}{d_1 d_2} \sum_{i,j} d_H^2(P_{i,j}, Q_{i,j}).$$

Finally, for an event $\mathcal{E}$,$\mathbb{1}_{[\mathcal{E}]}$ is the indicator function for that event, i.e., $\mathbb{1}_{[\mathcal{E}]}$ is 1 if $\mathcal{E}$ occurs and 0 otherwise.

1.6 Organization of the paper

We proceed in Section 2 by describing the 1-bit matrix completion problem in greater detail. In Section 3 we state our main results. Specifically, we propose a pair of convex programs for the 1-bit matrix completion problem and establish upper bounds on the accuracy with which these can recover the matrix $\boldsymbol{M}$ and the distribution of the observations $\boldsymbol{Y}$. We also establish lower bounds, showing that our upper bounds are nearly optimal. The proofs of these results are given in Section 4. Section 5 concludes with a brief discussion.

2 The 1-bit matrix completion problem

2.1 Observation model

We now introduce the more general observation model that we study in this paper. Given a matrix $\boldsymbol{M} \in \mathbb{R}^{d_1 \times d_2}$, a subset of indices $\Omega \subset [d_1] \times [d_2]$, and a differentiable function $f : \mathbb{R} \to [0,1]$, we observe

$$Y_{i,j} = \begin{cases} +1 & \text{with probability } f(M_{i,j}), \\ -1 & \text{with probability } 1 - f(M_{i,j}) \end{cases} \quad \text{for } (i,j) \in \Omega. \tag{2}$$

We will leave f general for now and discuss a few common choices just below. As has been important in previous work on matrix completion, we assume that Ω is chosen at random with $\mathbb{E}|\Omega| = m$. Specifically, we assume that Ω follows a binomial model in which each entry $(i,j) \in [d_1] \times [d_2]$ is included in Ω with probability $\frac{m}{d_1 d_2}$, independently.

Before discussing some particular choices for f, we first note that while the observation model described in (2) may appear on the surface to be somewhat different from the setup in (1), they are actually equivalent if f behaves likes a cumulative distribution function. Specifically, for the model in (1), if $\boldsymbol{Z}$ has i.i.d. entries, then by setting $f(x) := \mathbb{P}(Z_{1,1} \geq -x)$, the model in (1) reduces to that in (2). Similarly, for any choice of $f(x)$ in (2), if we define $\boldsymbol{Z}$ as having i.i.d. entries drawn from a distribution whose cumulative distribution function is given by $F_Z(x) = \mathbb{P}(z \leq x) = 1 - f(-x)$, then (2) reduces to (1). Of course, in any given situation one of these observation models may seem more or less natural than the other—for example, (1) may seem more appropriate when $\boldsymbol{M}$ is viewed as a latent variable which we might be interested in estimating, while (2) may make more sense when $\boldsymbol{M}$ is viewed as just a parameter of a distribution. Ultimately, however, the two models are equivalent.

We now consider two natural choices for f (or equivalently, for $\boldsymbol{Z}$):

Example 1 (Logistic regression/Logistic noise)**.** The logistic regression model, which is common in statistics, is captured by (2) with $f(x) = \frac{e^x}{1+e^x}$ and by (1) with $Z_{i,j}$ i.i.d. according to the standard logistic distribution.

Example 2 (Probit regression/Gaussian noise)**.** The probit regression model is captured by (2) by setting $f(x) = 1 - \Phi(-x/\sigma) = \Phi(x/\sigma)$ where Φ is the cumulative distribution function of a standard Gaussian and by (1) with $Z_{i,j}$ i.i.d. according to a mean-zero Gaussian distribution with variance σ^2.

2.2 Approximately low-rank matrices

The majority of the literature on matrix completion assumes that the first r singular values of $\boldsymbol{M}$ are nonzero and the remainder are exactly zero. However, in many applications the singular values instead exhibit only a gradual decay towards zero. Thus, in this paper we allow a relaxation of the assumption that $\boldsymbol{M}$ has rank exactly r. Instead, we assume that $\|\boldsymbol{M}\|_* \leq \alpha\sqrt{r d_1 d_2}$, where α is a parameter left to be determined, but which will often be of constant order. In other words, the singular values of $\boldsymbol{M}$ belong to a scaled ℓ_1 ball. In compressed sensing, belonging to an ℓ_p ball with $p \in (0,1]$ is a common relaxation of exact sparsity; in matrix completion, the nuclear-norm ball (or Schatten-1 ball) plays an analogous role.

The particular choice of scaling, $\alpha\sqrt{r d_1 d_2}$, arises from the following considerations. Suppose that each entry of $\boldsymbol{M}$ is bounded in magnitude by α and that $\operatorname{rank}(\boldsymbol{M}) \leq r$. Then

$$\|\boldsymbol{M}\|_* \leq \sqrt{r}\,\|\boldsymbol{M}\|_F \leq \sqrt{r d_1 d_2}\,\|\boldsymbol{M}\|_\infty \leq \alpha\sqrt{r d_1 d_2}.$$

Thus, the assumption that $\|\boldsymbol{M}\|_* \leq \alpha\sqrt{r d_1 d_2}$ is a relaxation of the conditions that $\operatorname{rank}(\boldsymbol{M}) \leq r$ and $\|\boldsymbol{M}\|_\infty \leq \alpha$. The condition that $\|\boldsymbol{M}\|_\infty \leq \alpha$ essentially means that the probability of seeing a $+1$ or -1 does not depend on the dimension. It is also a way of enforcing that $\boldsymbol{M}$ should not be too "spiky"; as discussed above this is an important assumption in order to make the recovery of $\boldsymbol{M}$ well-posed (e.g., see [36]).

3 Main results

We now state our main results. We will have two goals—the first is to accurately recover $\boldsymbol{M}$ itself, and the second is to accurately recover the distribution of $\boldsymbol{Y}$ given by $f(\boldsymbol{M})$.[1]

[1]Strictly speaking, $f(\boldsymbol{M}) \in [0,1]^{d_1 \times d_2}$ is simply a matrix of scalars, but these scalars implicitly define the distribution of $\boldsymbol{Y}$, so we will sometimes abuse notation slightly and refer to $f(\boldsymbol{M})$ as the distribution of $\boldsymbol{Y}$.

3.1 Convex programming

In order to approximate either $\boldsymbol{M}$ or $f(\boldsymbol{M})$, we will maximize the log-likelihood function of the optimization variable $\boldsymbol{X}$ given our observations subject to a set of convex constraints. In our case, the log-likelihood function is given by

$$F_{\Omega,\boldsymbol{Y}}(\boldsymbol{X}) := \sum_{(i,j)\in\Omega} \left(\mathbb{1}_{[Y_{i,j}=1]} \log(f(X_{i,j})) + \mathbb{1}_{[Y_{i,j}=-1]} \log(1 - f(X_{i,j})) \right).$$

To recover $\boldsymbol{M}$, we will use the solution to the following program:

$$\widehat{\boldsymbol{M}} = \arg\max_{\boldsymbol{X}} F_{\Omega,\boldsymbol{Y}}(\boldsymbol{X}) \qquad \text{subject to} \qquad \|\boldsymbol{X}\|_* \le \alpha\sqrt{rd_1d_2} \quad \text{and} \quad \|\boldsymbol{X}\|_\infty \le \alpha. \tag{3}$$

To recover the distribution $f(\boldsymbol{M})$, we need not enforce the infinity-norm constraint, and will use the following simpler program:

$$\widehat{\boldsymbol{M}} = \arg\max_{\boldsymbol{X}} F_{\Omega,\boldsymbol{Y}}(\boldsymbol{X}) \qquad \text{subject to} \qquad \|\boldsymbol{X}\|_* \le \alpha\sqrt{rd_1d_2} \tag{4}$$

In many cases, $F_{\Omega,\boldsymbol{Y}}(\boldsymbol{X})$ is a concave function and thus the above programs are convex. This can be easily checked in the case of the logistic model and can also be verified in the case of the probit model (e.g., see [48]).

3.2 Recovery of the matrix

We now state our main result concerning the recovery of the matrix $\boldsymbol{M}$. As discussed in Section 1.3 we place a "non-spikiness" condition on $\boldsymbol{M}$ to make recovery possible; we enforce this with an infinity-norm constraint. Further, some assumptions must be made on f for recovery of $\boldsymbol{M}$ to be feasible. We define two quantities L_α and β_α which control the "steepness" and "flatness" of f, respectively:

$$L_\alpha := \sup_{|x|\le\alpha} \frac{|f'(x)|}{f(x)(1-f(x))} \qquad \text{and} \qquad \beta_\alpha := \sup_{|x|\le\alpha} \frac{f(x)(1-f(x))}{(f'(x))^2}. \tag{5}$$

In this paper we will restrict our attention to f such that L_α and β_α are well-defined. In particular, we assume that f and f' are non-zero in $[-\alpha, \alpha]$. This assumption is fairly mild—for example, it includes the logistic and probit models (as we will see below in Remark 1). The quantity L_α appears only in our upper bounds, but it is generally well behaved. The quantity β_α appears both in our upper and lower bounds. Intuitively, it controls the "flatness" of f in the interval $[-\alpha, \alpha]$—the flatter f is, the larger β_α is. It is clear that some dependence on β_α is necessary. Indeed, if f is perfectly flat, then the magnitudes of the entries of $\boldsymbol{M}$ cannot be recovered, as seen in the noiseless case discussed in Section 1.3. Of course, when α is a fixed constant and f is a fixed function, both L_α and β_α are bounded by fixed constants independent of the dimension.

Theorem 1. *Assume that* $\|\boldsymbol{M}\|_* \le \alpha\sqrt{d_1d_2r}$ *and* $\|\boldsymbol{M}\|_\infty \le \alpha$. *Suppose that* Ω *is chosen at random following the binomial model of Section 2.1 with* $\mathbb{E}\,|\Omega| = m$. *Suppose that* $\boldsymbol{Y}$ *is generated as in* (2). *Let* L_α *and* β_α *be as in* (5). *Consider the solution* $\widehat{\boldsymbol{M}}$ *to* (3). *Then with probability at least* $1 - C_1/(d_1+d_2)$,

$$\frac{1}{d_1d_2}\|\widehat{\boldsymbol{M}} - \boldsymbol{M}\|_F^2 \le C_\alpha \sqrt{\frac{r(d_1+d_2)}{m}} \sqrt{1 + \frac{(d_1+d_2)\log(d_1d_2)}{m}}$$

with $C_\alpha := C_2\alpha L_\alpha\beta_\alpha$. *If* $m \ge (d_1+d_2)\log(d_1d_2)$ *then this simplifies to*

$$\frac{1}{d_1d_2}\|\widehat{\boldsymbol{M}} - \boldsymbol{M}\|_F^2 \le \sqrt{2}C_\alpha \sqrt{\frac{r(d_1+d_2)}{m}}. \tag{6}$$

Above, C_1 *and* C_2 *are absolute constants.*

Remark 1 (Recovery in the logistic and probit models). The logistic model satisfies the hypotheses of Theorem 1 with $\beta_\alpha = \frac{(1+e^\alpha)^2}{e^\alpha} \approx e^\alpha$ and $L_\alpha = 1$. The probit model has

$$\beta_\alpha \le c_1 \sigma^2 e^{\frac{\alpha^2}{2\sigma^2}} \quad \text{and} \quad L_\alpha \le c_2 \frac{\frac{\alpha}{\sigma}+1}{\sigma}$$

where we can take $c_1 = \pi$ and $c_2 = 8$. In particular, in the probit model the bound in (6) reduces to

$$\frac{1}{d_1 d_2} \|\widehat{\boldsymbol{M}} - \boldsymbol{M}\|_F^2 \le C \left(\frac{\alpha}{\sigma} + 1\right) \exp\left(\frac{\alpha^2}{2\sigma^2}\right) \sigma\alpha \sqrt{\frac{r(d_1+d_2)}{m}}. \tag{7}$$

Hence, when $\sigma < \alpha$, increasing the size of the noise leads to significantly improved error bounds—this is not an artifact of the proof. We will see in Section 3.4 that the exponential dependence on α in the logistic model (and on α/σ in the probit model) is intrinsic to the problem. Intuitively we should expect this since for such models, as $\|\boldsymbol{M}\|_\infty$ grows large, we essentially revert to the noiseless setting where estimation of $\boldsymbol{M}$ is impossible. Furthermore, in Section 3.4 we will also see that when α (or α/σ) is bounded by a constant, the error bound (6) is optimal up to a constant factor. Fortunately, in many applications, one would expect α to be small, and in particular to have little, if any, dependence on the dimension. This ensures that each measurement will always have a non-vanishing probability of returning 1 as well as a non-vanishing probability of returning -1.

Finally, note that if $\boldsymbol{M}$ is exactly rank r and satisfies $\|\boldsymbol{M}\|_\infty \le \alpha$, then as discussed in Section 2.2, $\boldsymbol{M}$ will automatically satisfy the assumptions of Theorem 1. Furthermore, note that the theorem also holds if $\Omega = [d_1] \times [d_2]$, i.e., if we sample each entry exactly once or observe a complete realization of $\boldsymbol{Y}$. Even in this context, the ability to recover $\boldsymbol{M}$ is somewhat surprising.

3.3 Recovery of the distribution

In many situations, we might not be interested in the underlying matrix $\boldsymbol{M}$, but rather in determining the distribution of the unknown entries of $\boldsymbol{Y}$. For example, in recommender systems, a natural question would be to determine the likelihood that a user would enjoy a particular unrated item.

Surprisingly, this distribution may be accurately recovered without any restriction on the infinity-norm of $\boldsymbol{M}$. This may be unexpected to those familiar with the matrix completion literature in which "non-spikiness" constraints seem to be unavoidable. In fact, we will show in Section 3.4 that the bound in Theorem 2 is near-optimal; further, we will show that even under the added constraint that $\|\boldsymbol{M}\|_\infty \le \alpha$, it would be impossible to estimate $f(\boldsymbol{M})$ significantly more accurately.

Theorem 2. *Assume that $\|\boldsymbol{M}\|_* \le \alpha\sqrt{d_1 d_2 r}$. Suppose that Ω is chosen at random following the binomial model of Section 2.1 with $\mathbb{E}|\Omega| = m$. Suppose that $\boldsymbol{Y}$ is generated as in (2), and let $L = \lim_{\alpha\to\infty} L_\alpha$. Let $\widehat{\boldsymbol{M}}$ be the solution to (4). Then, with probability at least $1 - C_1/(d_1+d_2)$,*

$$d_H^2(f(\widehat{\boldsymbol{M}}), f(\boldsymbol{M})) \le C_2 \alpha L \sqrt{\frac{r(d_1+d_2)}{m}} \sqrt{1 + \frac{(d_1+d_2)\log(d_1 d_2)}{m}}. \tag{8}$$

Furthermore, as long as $m \ge (d_1+d_2)\log(d_1 d_2)$, we have

$$d_H^2(f(\widehat{\boldsymbol{M}}), f(\boldsymbol{M})) \le \sqrt{2} C_2 \alpha L \sqrt{\frac{r(d_1+d_2)}{m}}. \tag{9}$$

Above, C_1 and C_2 are absolute constants.

While $L = 1$ for the logistic model, the astute reader will have noticed that for the probit model L is unbounded—that is, L_α tends to ∞ as $\alpha \to \infty$. L would also be unbounded for the case where $f(x)$ takes values of 1 or 0 outside of some range (as would be the case in (1) if the distribution of the noise had compact support). Fortunately, however, we can recover a result for these cases by enforcing an infinity-norm constraint, as described in Theorem 6 below. Moreover, for a large class of functions, f, L is indeed bounded. For example, in the latent variable version of (1) if the entries $Z_{i,j}$ are at least as fat-tailed as an exponential random variable, then L is bounded. To be more precise, suppose that f is continuously differentiable and for simplicity assume that the distribution of $Z_{i,j}$ is symmetric and $|f'(x)|/(1-f(x))$ is monotonic for x sufficiently large. If $\mathbb{P}(|Z_{i,j}| \geq t) \geq C\exp(-ct)$ for all $t \geq 0$, then one can show that L is finite. This property is also essentially equivalent to the requirement that a distribution have bounded *hazard rate*. As noted above, this property holds for the logistic distribution, but also for many other common distributions, including the Laplacian, student's t, Cauchy, and others.

3.4 Room for improvement?

We now discuss the extent to which Theorems 1 and 2 are optimal. We give three theorems, all proved using information theoretic methods, which show that these results are nearly tight, even when some of our assumptions are relaxed. Theorem 3 gives a lower bound to nearly match the upper bound on the error in recovering $\boldsymbol{M}$ derived in Theorem 1. Theorem 4 compares our upper bounds to those available without discretization and shows that very little is lost when discretizing to a single bit. Finally, Theorem 5 gives a lower bound matching, up to a constant factor, the upper bound on the error in recovering the distribution $f(\boldsymbol{M})$ given in Theorem 2. Theorem 5 also shows that Theorem 2 does not suffer by dropping the canonical "spikiness" constraint.

Our lower bounds require a few assumptions, so before we delve into the bounds themselves, we briefly argue that these assumptions are rather innocuous. First, without loss of generality (since we can always adjust f to account for rescaling $\boldsymbol{M}$), we assume that $\alpha \geq 1$. Next, we require that the parameters be sufficiently large so that

$$\alpha^2 r \max\{d_1, d_2\} \geq C_0 \tag{10}$$

for an absolute constant C_0. Note that we could replace this with a simpler, but still mild, condition that $d_1 > C_0$. Finally, we also require that $r \geq c$ where c is either 1 or 4 and that $r \leq O(\min\{d_1, d_2\}/\alpha^2)$, where $O(\cdot)$ hides parameters (which may differ in each Theorem) that we make explicit below. This last assumption simply means that we are in the situation where r is significantly smaller than d_1 and d_2, i.e., the matrix is of approximately low rank.

In the following, let

$$K = \left\{\boldsymbol{M} \ : \ \|\boldsymbol{M}\|_* \leq \alpha\sqrt{rd_1d_2}, \|\boldsymbol{M}\|_\infty \leq \alpha\right\} \tag{11}$$

denote the set of matrices whose recovery is guaranteed by Theorem 1.

3.4.1 Recovery from 1-bit measurements

Theorem 3. *Fix α, r, d_1, and d_2 to be such that $r \geq 4$ and (10) holds. Let β_α be defined as in (5), and suppose that $f'(x)$ is decreasing for $x > 0$. Let Ω be any subset of $[d_1] \times [d_2]$ with cardinality m, and let $\boldsymbol{Y}$ be as in (2). Consider any algorithm which, for any $\boldsymbol{M} \in K$, takes as input $Y_{i,j}$ for $(i,j) \in \Omega$ and returns $\widehat{\boldsymbol{M}}$. Then, there exists $\boldsymbol{M} \in K$ such that with probability at least $3/4$,*

$$\frac{1}{d_1 d_2}\|\boldsymbol{M} - \widehat{\boldsymbol{M}}\|_F^2 \geq \min\left\{C_1, C_2\alpha\sqrt{\beta_{\frac{3}{4}\alpha}}\sqrt{\frac{r\max\{d_1,d_2\}}{m}}\right\} \tag{12}$$

as long as the right-hand side of (12) exceeds $r\alpha^2/\min(d_1, d_2)$. Above, C_1 and C_2 are absolute constants.[2]

[2]Here and in the theorems below, the choice of $3/4$ in the probability bound is arbitrary, and can be adjusted at the cost of changing C_0 in (10) and C_1 and C_2. Similarly, $\beta_{\frac{3}{4}\alpha}$ can be replaced by $\beta_{(1-\epsilon)\alpha}$ for any $\epsilon > 0$.

The requirement that the right-hand side of (12) be larger than $r\alpha^2/\min(d_1, d_2)$ is satisfied as long as $r \leq O(\min\{d_1, d_2\}/\alpha^2)$. In particular, it is satisfied whenever

$$r \leq C_3 \frac{\min(1, \beta_0) \cdot \min(d_1, d_2)}{\alpha^2}$$

for a fixed constant C_3. Note also that in the latent variable model in (1), $f'(x)$ is simply the probability density of $Z_{i,j}$. Thus, the requirement that $f'(x)$ be decreasing is simply asking the probability density to have decreasing tails. One can easily check that this is satisfied for the logistic and probit models.

Note that if α is bounded by a constant and f is fixed (in which case β_α and $\beta_{\alpha'}$ are bounded by a constant), then the lower bound of Theorem 3 matches the upper bound given in (6) up to a constant. When α is not treated as a constant, the bounds differ by a factor of $\sqrt{\beta_\alpha}$. In the logistic model $\beta_\alpha \approx e^\alpha$ and so this amounts to the difference between $e^{\alpha/2}$ and e^α. The probit model has a similar change in the constant of the exponent.

3.4.2 Recovery from unquantized measurements

Next we show that, surprisingly, very little is lost by discretizing to a single bit. In Theorem 4, we consider an "unquantized" version of the latent variable model in (1) with Gaussian noise. That is, let $\boldsymbol{Z}$ be a matrix of i.i.d. Gaussian random variables, and suppose the noisy entries $M_{i,j} + Z_{i,j}$ are observed directly, without discretization. In this setting, we give a lower bound that still nearly matches the upper bound given in Theorem 1, up to the β_α term.

Theorem 4. *Fix α, r, d_1, and d_2 to be such that $r \geq 1$ and (10) holds. Let Ω be any subset of $[d_1] \times [d_2]$ with cardinality m, and let $\boldsymbol{Z}$ be a $d_1 \times d_2$ matrix with i.i.d. Gaussian entries with variance σ^2. Consider any algorithm which, for any $\boldsymbol{M} \in K$, takes as input $Y_{i,j} = M_{i,j} + Z_{i,j}$ for $(i, j) \in \Omega$ and returns $\widehat{\boldsymbol{M}}$. Then, there exists $\boldsymbol{M} \in K$ such that with probability at least $3/4$,*

$$\frac{1}{d_1 d_2} \|\boldsymbol{M} - \widehat{\boldsymbol{M}}\|_F^2 \geq \min\left\{C_1, C_2 \alpha\sigma \sqrt{\frac{r \max\{d_1, d_2\}}{m}}\right\} \tag{13}$$

as long as the right-hand side of (13) exceeds $r\alpha^2/\min(d_1, d_2)$. Above, C_1 and C_2 are absolute constants.

The requirement that the right-hand side of (13) be larger than $r\alpha^2/\min(d_1, d_2)$ is satisfied whenever

$$r \leq C_3 \frac{\min(1, \sigma^2) \min(d_1, d_2)}{\alpha^2}$$

for a fixed constant C_3.

Following Remark 1, the lower bound given in (13) matches the upper bound proven in Theorem 1 for the solution to (4) up to a constant, as long as α/σ is bounded by a constant. In other words:

When the signal-to-noise ratio is constant, almost nothing is lost by quantizing to a single bit.

Perhaps it is not particularly surprising that 1-bit quantization induces little loss of information in the regime where the noise is comparable to the underlying quantity we wish to estimate—however, what *is* somewhat of a surprise is that the simple convex program in (4) can successfully recover all of the information contained in these 1-bit measurements.

Before proceeding, we also briefly note that our Theorem 4 is somewhat similar to Theorem 3 in [36]. The authors in [36] consider slightly different sets K: these sets are more restrictive in the sense that it is required that $\alpha \geq \sqrt{32 \log n}$ and less restrictive because the nuclear-norm constraint may be replaced by a general Schatten-p norm constraint. It was important for us to allow $\alpha = O(1)$ in order to compare with our upper bounds due to the exponential dependence of β_α on α in Theorem 1 for the probit model. This led to some new challenges in the proof. Finally, it is also noteworthy that our statements hold for arbitrary sets Ω, while the argument in [36] is only valid for a random choice of Ω.

3.4.3 Recovery of the distribution from 1-bit measurements

To conclude we address the optimality of Theorem 2. We show that under mild conditions on f, any algorithm that recovers the distribution $f(\boldsymbol{M})$ must yield an estimate whose Hellinger distance deviates from the true distribution by an amount proportional to $\alpha\sqrt{rd_1d_2/m}$, matching the upper bound of (9) up to a constant. Notice that the lower bound holds even if the algorithm is promised that $\|\boldsymbol{M}\|_\infty \leq \alpha$, which the upper bound did not require.

Theorem 5. *Fix α, r, d_1, and d_2 to be such that $r \geq 4$ and (10) holds. Let L_1 be defined as in (5), and suppose that $f'(x) \geq c$ and $c' \leq f(x) \leq 1 - c'$ for $x \in [-1, 1]$, for some constants $c, c' > 0$. Let Ω be any subset of $[d_1] \times [d_2]$ with cardinality m, and let $\boldsymbol{Y}$ be as in (2). Consider any algorithm which, for any $\boldsymbol{M} \in K$, takes as input $Y_{i,j}$ for $(i,j) \in \Omega$ and returns $\widehat{\boldsymbol{M}}$. Then, there exists $\boldsymbol{M} \in K$ such that with probability at least $3/4$,*

$$d_H^2(f(\boldsymbol{M}), f(\widehat{\boldsymbol{M}})) \geq \min\left\{C_1, C_2 \frac{\alpha}{L_1}\sqrt{\frac{r\max\{d_1,d_2\}}{m}}\right\} \tag{14}$$

as long as the right-hand side of (14) exceeds $r\alpha^2/\min(d_1,d_2)$. Above, C_1 and C_2 are constants that depend on c, c'.

The requirement that the right-hand side of (14) be larger than $r\alpha^2/\min(d_1,d_2)$ is satisfied whenever

$$r \leq C_3 \min\left(1, \frac{1}{L_1^2}\right)\frac{\min(d_1,d_2)}{\alpha^2}$$

for a constant C_3 that depends only on c, c'. Note also that the condition that f and f' be well-behaved in the interval $[-1,1]$ is satisfied for the logistic model with $c = 1/4$ and $c' = \frac{1}{1+e} \leq 0.269$. Similarly, we may take $c = 0.242$ and $c' = 0.159$ in the probit model.

4 Proofs of the main results

In this section we provide the proofs of the main theorems presented in Section 3. To begin, we first define some additional notation that we will need for the proofs. For two probability distributions $\mathcal{P}$ and $\mathcal{Q}$ on a finite set A, $D(\mathcal{P}\|\mathcal{Q})$ will denote the Kullback-Leibler (KL) divergence,

$$D(\mathcal{P}\|\mathcal{Q}) = \sum_{x\in A}\mathcal{P}(x)\log\left(\frac{\mathcal{P}(x)}{\mathcal{Q}(x)}\right),$$

where $\mathcal{P}(x)$ denotes the probability of the outcome x under the distribution $\mathcal{P}$. We will abuse this notation slightly by overloading it in two ways. First, for scalar inputs $p, q \in [0,1]$, we will set

$$D(p\|q) = p\log\left(\frac{p}{q}\right) + (1-p)\log\left(\frac{1-p}{1-q}\right).$$

Second, for two matrices $\boldsymbol{P}, \boldsymbol{Q} \in [0,1]^{d_1\times d_2}$, we define

$$D(\boldsymbol{P}\|\boldsymbol{Q}) = \frac{1}{d_1d_2}\sum_{i,j} D(P_{i,j}\|Q_{i,j}).$$

We first prove Theorem 2. Theorem 1 will then follow from an approximation argument. Finally, our lower bounds will be proved in Section 4.3 using information theoretic arguments.

4.1 Proof of Theorem 2

We will actually prove a slightly more general statement, which will be helpful in the proof of Theorem 1. We will assume that $\|\boldsymbol{M}\|_\infty \leq \gamma$, and we will modify the program (4) to enforce $\|\boldsymbol{X}\|_\infty \leq \gamma$. That is, we will consider the program

$$\widehat{\boldsymbol{M}} = \arg\max_{\boldsymbol{X}} F_{\Omega,\boldsymbol{Y}}(\boldsymbol{X}) \qquad \text{subject to} \qquad \|\boldsymbol{X}\|_* \leq \alpha\sqrt{rd_1d_2} \quad \text{and} \quad \|\boldsymbol{X}\|_\infty \leq \gamma. \tag{15}$$

We will then send $\gamma \to \infty$ to recover the statement of Theorem 2. Formally, we prove the following theorem.

Theorem 6. *Assume that $\|\boldsymbol{M}\|_* \leq \alpha\sqrt{rd_1d_2}$ and $\|\boldsymbol{M}\|_\infty \leq \gamma$. Suppose that Ω is chosen at random following the binomial model of Section 2.1 and satisfying $\mathbb{E}|\Omega| = m$. Suppose that $\boldsymbol{Y}$ is generated as in (2), and let L_γ be as in (5). Let $\widehat{\boldsymbol{M}}$ be the solution to (15). Then, with probability at least $1 - C_1/(d_1+d_2)$,*

$$d_H^2(f(\widehat{\boldsymbol{M}}), f(\boldsymbol{M})) \leq C_2 L_\gamma \alpha \sqrt{\frac{r(d_1+d_2)}{m}}\sqrt{1 + \frac{(d_1+d_2)\log(d_1d_2)}{m}}. \tag{16}$$

Above, C_1 and C_2 are absolute constants.

The key to proving Theorem 6 will be to establish the following concentration inequality.

Lemma 1. *Let $G \subset \mathbb{R}^{d_1 \times d_2}$ be*

$$G = \left\{ \boldsymbol{X} \in \mathbb{R}^{d_1 \times d_2} \ : \ \|\boldsymbol{X}\|_* \leq \alpha\sqrt{rd_1d_2} \right\}$$

for some $r \leq \min\{d_1, d_2\}$ and $\alpha \geq 0$. Then

$$\mathbb{P}\left(\sup_{\boldsymbol{X}\in G} |F_{\Omega,\boldsymbol{Y}}(\boldsymbol{X}) - \mathbb{E}\, F_{\Omega,\boldsymbol{Y}}(\boldsymbol{X})| \geq C_0 \alpha L_\gamma \sqrt{r}\sqrt{m(d_1+d_2) + d_1d_2\log(d_1d_2)} \right) \leq \frac{C_1}{d_1+d_2}, \tag{17}$$

where C_0 and C_1 are absolute constants and the probability and the expectation are both over the choice of Ω and the draw of Y.

We will prove this lemma below, but first we show how it implies Theorem 6. To begin, notice that for any choice of $\boldsymbol{X}$,

$$\begin{aligned}\mathbb{E}\left[F_{\Omega,\boldsymbol{Y}}(\boldsymbol{X}) - F_{\Omega,\boldsymbol{Y}}(\boldsymbol{M})\right] &= \frac{m}{d_1d_2}\sum_{i,j}\left(f(M_{i,j})\log\left(\frac{f(X_{i,j})}{f(M_{i,j})}\right) + (1 - f(M_{i,j}))\log\left(\frac{1-f(X_{i,j})}{1-f(M_{i,j})}\right)\right) \\ &= -mD(f(\boldsymbol{M})\|f(\boldsymbol{X})),\end{aligned}$$

where the expectation is over both Ω and $\boldsymbol{Y}$. Next, note that by assumption $\boldsymbol{M} \in G$. Moreover, from the definition of $\widehat{\boldsymbol{M}}$ we also have that $\widehat{\boldsymbol{M}} \in G$ and $F_{\Omega,\boldsymbol{Y}}(\widehat{\boldsymbol{M}}) \geq F_{\Omega,\boldsymbol{Y}}(\boldsymbol{M})$. Thus, we can write

$$\begin{aligned}0 &\leq F_{\Omega,\boldsymbol{Y}}(\widehat{\boldsymbol{M}}) - F_{\Omega,\boldsymbol{Y}}(\boldsymbol{M}) \\ &= F_{\Omega,\boldsymbol{Y}}(\widehat{\boldsymbol{M}}) + \mathbb{E}\, F_{\Omega,\boldsymbol{Y}}(\widehat{\boldsymbol{M}}) - \mathbb{E}\, F_{\Omega,\boldsymbol{Y}}(\widehat{\boldsymbol{M}}) + \mathbb{E}\, F_{\Omega,\boldsymbol{Y}}(\boldsymbol{M}) - \mathbb{E}\, F_{\Omega,\boldsymbol{Y}}(\boldsymbol{M}) - F_{\Omega,\boldsymbol{Y}}(\boldsymbol{M}) \\ &\leq \mathbb{E}\left[F_{\Omega,\boldsymbol{Y}}(\widehat{\boldsymbol{M}}) - F_{\Omega,\boldsymbol{Y}}(\boldsymbol{M})\right] + |F_{\Omega,\boldsymbol{Y}}(\widehat{\boldsymbol{M}}) - \mathbb{E}\, F_{\Omega,\boldsymbol{Y}}(\widehat{\boldsymbol{M}})| + |F_{\Omega,\boldsymbol{Y}}(\boldsymbol{M}) - \mathbb{E}\, F_{\Omega,\boldsymbol{Y}}(\boldsymbol{M})| \\ &\leq -mD(f(\boldsymbol{M})\|f(\widehat{\boldsymbol{M}})) + 2\sup_{\boldsymbol{X}\in G}|F_{\Omega,\boldsymbol{Y}}(\boldsymbol{X}) - \mathbb{E}\, F_{\Omega,\boldsymbol{Y}}(\boldsymbol{X})|.\end{aligned}$$

Applying Lemma 1, we obtain that with probability at least $1 - C_1/(d_1+d_2)$, we have

$$0 \leq -mD(f(\boldsymbol{M})\|f(\widehat{\boldsymbol{M}})) + 2C_0\alpha L_\gamma\sqrt{r}\sqrt{m(d_1+d_2) + d_1d_2\log(d_1d_2)}$$

In this case, by rearranging and applying the fact that $\sqrt{d_1 d_2} \le d_1 + d_2$, we obtain

$$D(f(\boldsymbol{M}) \| f(\widehat{\boldsymbol{M}})) \le 2C_0 \alpha L_\gamma \sqrt{\frac{r(d_1+d_2)}{m}} \sqrt{1 + \frac{(d_1+d_2)\log(d_1 d_2)}{m}} \tag{18}$$

Finally, we note that the KL divergence can easily be bounded below by the Hellinger distance:

$$d_H^2(p,q) \le D(p\|q).$$

This is a simple consequence of Jensen's inequality combined with the fact that $1 - x \le -\log x$. Thus, from (18) we obtain

$$d_H^2(f(\boldsymbol{M}), f(\widehat{\boldsymbol{M}})) \le 2C_0 \alpha L_\gamma \sqrt{\frac{r(d_1+d_2)}{m}} \sqrt{1 + \frac{(d_1+d_2)\log(d_1 d_2)}{m}},$$

which establishes Theorem 6. Theorem 2 then follows by taking the limit as $\gamma \to \infty$.

Proof of Lemma 1. We begin by noting that for any $h > 0$, by using Markov's inequality we have that

$$\begin{aligned}
&\mathbb{P}\left(\sup_{\boldsymbol{X} \in G} |F_{\Omega,\boldsymbol{Y}}(\boldsymbol{X}) - \mathbb{E}\, F_{\Omega,\boldsymbol{Y}}(\boldsymbol{X})| \ge C_0 \alpha L_\gamma \sqrt{r}\sqrt{m(d_1+d_2) + d_1 d_2 \log(d_1 d_2)}\right) \\
&= \mathbb{P}\left(\sup_{\boldsymbol{X} \in G} |F_{\Omega,\boldsymbol{Y}}(\boldsymbol{X}) - \mathbb{E}\, F_{\Omega,\boldsymbol{Y}}(\boldsymbol{X})|^h \ge \left(C_0 \alpha L_\gamma \sqrt{r}\sqrt{m(d_1+d_2) + d_1 d_2 \log(d_1 d_2)}\right)^h\right) \\
&\le \frac{\mathbb{E}\left[\sup_{\boldsymbol{X} \in G} |F_{\Omega,\boldsymbol{Y}}(\boldsymbol{X}) - \mathbb{E}\, F_{\Omega,\boldsymbol{Y}}(\boldsymbol{X})|^h\right]}{\left(C_0 \alpha L_\gamma \sqrt{r}\sqrt{m(d_1+d_2) + d_1 d_2 \log(d_1 d_2)}\right)^h}. \qquad (19)
\end{aligned}$$

The bound in (17) will follow by combining this with an upper bound on $\mathbb{E}\left[\sup_{\boldsymbol{X} \in G} |F_{\Omega,\boldsymbol{Y}}(\boldsymbol{X}) - \mathbb{E}\, F_{\Omega,\boldsymbol{Y}}(\boldsymbol{X})|^h\right]$ and setting $h = \log(d_1 + d_2)$. Towards this end, note that we can write the definition of $F_{\Omega,\boldsymbol{Y}}$ as

$$F_{\Omega,\boldsymbol{Y}}(\boldsymbol{X}) = \sum_{i,j} \left(\mathbb{1}_{[(i,j)\in\Omega]} \left(\mathbb{1}_{[Y_{i,j}=1]} \log(f(X_{i,j})) + \mathbb{1}_{[Y_{i,j}=-1]} \log(1 - f(X_{i,j}))\right)\right).$$

By a symmetrization argument (Lemma 6.3 in [31]),

$$\begin{aligned}
&\mathbb{E}\left[\sup_{\boldsymbol{X} \in G} |F_{\Omega,\boldsymbol{Y}}(\boldsymbol{X}) - \mathbb{E}\, F_{\Omega,\boldsymbol{Y}}(\boldsymbol{X})|^h\right] \\
&\qquad \le 2^h\, \mathbb{E}\left[\sup_{\boldsymbol{X} \in G} \left|\sum_{i,j} \varepsilon_{i,j} \mathbb{1}_{[(i,j)\in\Omega]} \left(\mathbb{1}_{[Y_{i,j}=1]} \log(f(X_{i,j})) + \mathbb{1}_{[Y_{i,j}=-1]} \log(1 - f(X_{i,j}))\right)\right|^h\right],
\end{aligned}$$

where the $\varepsilon_{i,j}$ are i.i.d. Rademacher random variables and the expectation in the upper bound is with respect to both Ω and $\boldsymbol{Y}$ as well as with respect to the $\varepsilon_{i,j}$. To bound the latter term, we apply a contraction principle (Theorem 4.12 in [31]). By the definition of L_γ and the assumption that $\|\widehat{\boldsymbol{M}}\|_\infty \le \gamma$, both $\log(f(x))/L_\gamma$ and $\log(1 - f(x))/L_\gamma$ are contractions. Thus, up to a factor of 2, the expected value of the supremum can only decrease when $\log(f(X_{i,j}))$ is replaced by $X_{i,j}$ and similarly $\log(1 - f(X_{i,j}))$ by $-X_{i,j}$. Thus we obtain

$$\begin{aligned}
\mathbb{E}\left[\sup_{\boldsymbol{X} \in G} |F_{\Omega,\boldsymbol{Y}}(\boldsymbol{X}) - \mathbb{E}\, F_{\Omega,\boldsymbol{Y}}(\boldsymbol{X})|^h\right] &\le 2^h (2L_\gamma)^h\, \mathbb{E}\left[\sup_{\boldsymbol{X} \in G} \left|\sum_{i,j} \varepsilon_{i,j} \mathbb{1}_{[(i,j)\in\Omega]} \left(\mathbb{1}_{[Y_{i,j}=1]} X_{i,j} - \mathbb{1}_{[Y_{i,j}=-1]} X_{i,j}\right)\right|^h\right] \\
&= (4L_\gamma)^h\, \mathbb{E}\left[\sup_{\boldsymbol{X} \in G} |\langle \Delta_\Omega \circ \boldsymbol{E} \circ \boldsymbol{Y}, \boldsymbol{X}\rangle|^h\right], \qquad (20)
\end{aligned}$$

where $\boldsymbol{E}$ denotes the matrix with entries given by $\varepsilon_{i,j}$, Δ_Ω denotes the indicator matrix for Ω (so that $[\Delta_\Omega]_{i,j} = 1$ if $(i,j) \in \Omega$ and 0 otherwise), and $\circ$ denotes the Hadamard product. Using the facts that the distribution of $\boldsymbol{E} \circ \boldsymbol{Y}$ is the same as the distribution of $\boldsymbol{E}$ and that $|\langle \boldsymbol{A}, \boldsymbol{B} \rangle| \le \|\boldsymbol{A}\| \|\boldsymbol{B}\|_*$, we have that

$$\begin{aligned}
\mathbb{E}\left[\sup_{\boldsymbol{X}\in G} |\langle \Delta_\Omega \circ \boldsymbol{E} \circ \boldsymbol{Y}, \boldsymbol{X}\rangle|^h\right] &= \mathbb{E}\left[\sup_{\boldsymbol{X}\in G} |\langle \boldsymbol{E} \circ \Delta_\Omega, \boldsymbol{X}\rangle|^h\right] \\
&\le \mathbb{E}\left[\sup_{\boldsymbol{X}\in G} \|\boldsymbol{E} \circ \Delta_\Omega\|^h \|\boldsymbol{X}\|_*^h\right] \\
&= \left(\alpha\sqrt{d_1 d_2 r}\right)^h \mathbb{E}\left[\|\boldsymbol{E} \circ \Delta_\Omega\|^h\right], \qquad (21)
\end{aligned}$$

To bound $\mathbb{E}\left[\|\boldsymbol{E} \circ \Delta_\Omega\|^h\right]$, observe that $\boldsymbol{E} \circ \Delta_\Omega$ is a matrix with i.i.d. zero mean entries and thus by Theorem 1.1 of [43],

$$\mathbb{E}\left[\|\boldsymbol{E} \circ \Delta_\Omega\|^h\right] \le C\left(\mathbb{E}\left[\max_{1\le i\le d_1}\left(\sum_{j=1}^{d_2}\Delta_{i,j}\right)^{\frac{h}{2}}\right] + \mathbb{E}\left[\max_{1\le j\le d_2}\left(\sum_{i=1}^{d_1}\Delta_{i,j}\right)^{\frac{h}{2}}\right]\right)$$

for some constant C. This in turn implies that

$$\left(\mathbb{E}\left[\|\boldsymbol{E} \circ \Delta_\Omega\|^h\right]\right)^{\frac{1}{h}} \le C^{\frac{1}{h}}\left(\left(\mathbb{E}\left[\max_{1\le i\le d_1}\left(\sum_{j=1}^{d_2}\Delta_{i,j}\right)^{\frac{h}{2}}\right]\right)^{\frac{1}{h}} + \left(\mathbb{E}\left[\max_{1\le j\le d_2}\left(\sum_{i=1}^{d_1}\Delta_{i,j}\right)^{\frac{h}{2}}\right]\right)^{\frac{1}{h}}\right). \qquad (22)$$

We first focus on the row sum $\sum_{j=1}^{d_2}\Delta_{i,j}$ for a particular choice of i. Using Bernstein's inequality, for all $t > 0$ we have

$$\mathbb{P}\left(\left|\sum_{j=1}^{d_2}\left(\Delta_{i,j} - \frac{m}{d_1 d_2}\right)\right| > t\right) \le 2\exp\left(\frac{-t^2/2}{m/d_1 + t/3}\right).$$

In particular, if we set $t \ge 6m/d_1$, then for each i we have

$$\mathbb{P}\left(\left|\sum_{j=1}^{d_2}\left(\Delta_{i,j} - \frac{m}{d_1 d_2}\right)\right| > t\right) \le 2\exp(-t) = 2\mathbb{P}\left(W_i > t\right), \qquad (23)$$

where $W_1, \ldots, W_{d_1}$ are i.i.d. exponential random variables.

Below we use the fact that for any positive random variable q we can write $\mathbb{E}\, q = \int_0^\infty \mathbb{P}\left(q \ge t\right)$, allowing us to bound

$$\begin{aligned}
\left(\mathbb{E}\left[\max_{1\le i\le d_1}\left(\sum_{j=1}^{d_2}\Delta_{i,j}\right)^{\frac{h}{2}}\right]\right)^{\frac{1}{h}} &\le \sqrt{\frac{m}{d_1}} + \left(\mathbb{E}\left[\max_{1\le i\le d_1}\left|\sum_{j=1}^{d_2}\left(\Delta_{i,j} - \frac{m}{d_1 d_2}\right)\right|^{\frac{h}{2}}\right]\right)^{\frac{1}{h}} \\
&\le \sqrt{\frac{m}{d_1}} + \left(\mathbb{E}\left[\max_{1\le i\le d_1}\left|\sum_{j=1}^{d_2}\left(\Delta_{i,j} - \frac{m}{d_1 d_2}\right)\right|^{h}\right]\right)^{\frac{1}{2h}} \\
&= \sqrt{\frac{m}{d_1}} + \left(\int_0^\infty \mathbb{P}\left(\max_{1\le i\le d_1}\left|\sum_{j=1}^{d_2}\left(\Delta_{i,j} - \frac{m}{d_1 d_2}\right)\right|^{h} \ge t\right) dt\right)^{\frac{1}{2h}}
\end{aligned}$$

$$
\begin{aligned}
&\le \sqrt{\frac{m}{d_1}} + \left(\left(\frac{6m}{d_1}\right)^h + \int_{(6m/d_1)^h}^{\infty} \mathbb{P}\left(\max_{1\le i\le d_1} \left| \sum_{j=1}^{d_2} \left(\Delta_{i,j} - \frac{m}{d_1 d_2}\right) \right|^h \ge t \right) dt \right)^{\frac{1}{2h}} \\
&\le \sqrt{\frac{m}{d_1}} + \left(\left(\frac{6m}{d_1}\right)^h + 2\int_{(6m/d_1)^h}^{\infty} \mathbb{P}\left(\max_{1\le i\le d_1} W_i^h \ge t \right) dt \right)^{\frac{1}{2h}} \\
&\le \sqrt{\frac{m}{d_1}} + \left(\left(\frac{6m}{d_1}\right)^h + 2\mathbb{E}\left[(\max_{1\le i\le d_1} W_i)^h \right] \right)^{\frac{1}{2h}}.
\end{aligned}
$$

Above, we have used the triangle inequality in the first line, followed by Jensen's inequality in the second line. In the fifth line, (23), along with independence, allows us to introduce $\max_i W_i$. By standard computations for exponential random variables,

$$
\begin{aligned}
\mathbb{E}\left[\max_{1\le i\le d_1} W_i^h \right] &\le \mathbb{E}\left[\left| \max_{1\le i\le d_1} W_i - \log d_1 \right|^h \right] + \log^h(d_1) \\
&\le 2h! + \log^h(d_1).
\end{aligned}
$$

Thus, we have

$$
\begin{aligned}
\left(\mathbb{E}\left[\max_{1\le i\le d_1} \left(\sum_{j=1}^{d_2} \Delta_{i,j} \right)^{\frac{h}{2}} \right] \right)^{\frac{1}{h}} &\le \sqrt{\frac{m}{d_1}} + \left(\left(\frac{6m}{d_1}\right)^h + 2\left(\log^h(d_1) + 2(h!)\right) \right)^{\frac{1}{2h}} \\
&\le (1+\sqrt{6})\sqrt{\frac{m}{d_1}} + 2^{\frac{1}{2h}}\left(\sqrt{\log d_1} + 2^{\frac{1}{2h}}\sqrt{h} \right) \\
&\le (1+\sqrt{6})\sqrt{\frac{m}{d_1}} + (2+\sqrt{2})\sqrt{\log(d_1+d_2)},
\end{aligned}
$$

using the choice $h = \log(d_1+d_2) \ge 1$ in the final line.

A similar argument bounds the column sums, and thus from (22) we conclude that

$$
\begin{aligned}
\left(\mathbb{E}\left[\|\boldsymbol{E} \circ \Delta_\Omega\|^h \right] \right)^{\frac{1}{h}} &\le C^{\frac{1}{h}} \left((1+\sqrt{6})\left(\sqrt{\frac{m}{d_1}} + \sqrt{\frac{m}{d_2}} \right) + (2+\sqrt{2})\sqrt{\log(d_1+d_2)} \right) \\
&\le C^{\frac{1}{h}} \left((1+\sqrt{6})\left(\sqrt{\frac{2m(d_1+d_2)}{d_1 d_2}} \right) + (2+\sqrt{2})\sqrt{\log(d_1+d_2)} \right) \\
&\le C^{\frac{1}{h}} 2(1+\sqrt{6}) \sqrt{\frac{m(d_1+d_2) + d_1 d_2 \log(d_1+d_2)}{d_1 d_2}},
\end{aligned}
$$

where the second and third inequalities both follow from Jensen's inequality. Combining this with (20) and (21), we obtain

$$
\left(\mathbb{E}\left[\sup_{\boldsymbol{X}\in G} |F_{\Omega,\boldsymbol{Y}}(\boldsymbol{X}) - \mathbb{E}\, F_{\Omega,\boldsymbol{Y}}(\boldsymbol{X})|^h \right] \right)^{\frac{1}{h}} \le C^{\frac{1}{h}} 8(1+\sqrt{6})\alpha L_\gamma \sqrt{r}\sqrt{m(d_1+d_2) + d_1 d_2 \log(d_1+d_2)}.
$$

Plugging this into 19 we obtain that the probability in 19 is upper bounded by

$$
C\left(\frac{8(1+\sqrt{6})}{C_0} \right)^{\log(d_1+d_2)} \le \frac{C}{d_1+d_2},
$$

provided that $C_0 \ge 8(1+\sqrt{6})/e$, which establishes the lemma. □

4.2 Proof of Theorem 1

The proof of Theorem 1 follows immediately from Theorem 6 (with $\gamma = \alpha$) combined with the following lemma.

Lemma 2. *Let f be a differentiable function and let $\|\boldsymbol{M}\|_\infty, \|\widehat{\boldsymbol{M}}\|_\infty \le \alpha$. Then*

$$d_H^2(f(\boldsymbol{M}), f(\widehat{\boldsymbol{M}})) \ge \inf_{|\xi|\le\alpha} \frac{(f'(\xi))^2}{8f(\xi)(1-f(\xi))} \frac{\|\boldsymbol{M} - \widehat{\boldsymbol{M}}\|_F^2}{d_1 d_2}.$$

Proof. For any pair of entries $x = M_{i,j}$ and $y = \widehat{M}_{i,j}$, write

$$\left(\sqrt{f(x)} - \sqrt{f(y)}\right)^2 + \left(\sqrt{1-f(x)} - \sqrt{1-f(y)}\right)^2 \ge \frac{1}{2}\left(\left(\sqrt{f(x)} - \sqrt{f(y)}\right) - \left(\sqrt{f(x)} - \sqrt{f(y)}\right)\right)^2.$$

Using Taylor's theorem to expand the quantity inside the square, for some ξ between x and y,

$$\begin{aligned}\left(\sqrt{f(x)} - \sqrt{f(y)}\right)^2 + \left(\sqrt{1-f(x)} - \sqrt{1-f(y)}\right)^2 &\ge \frac{1}{2}\left(\frac{f'(\xi)(y-x)}{2\sqrt{f(\xi)}} + \frac{f'(\xi)(y-x)}{2\sqrt{1-f(\xi)}}\right)^2 \\ &\ge \frac{1}{8}(f'(\xi))^2(y-x)^2\left(\frac{1}{f(\xi)} + \frac{1}{1-f(\xi)}\right) \\ &= \frac{(f'(\xi))^2}{8f(\xi)(1-f(\xi))}(y-x)^2.\end{aligned}$$

The lemma follows by summing across all entries and dividing by $d_1 d_2$. ☐

4.3 Lower bounds

The proofs of our lower bounds each follow a similar outline, using classical information theoretic techniques that have also proven useful in the context of compressed sensing [8, 39]. At a high level, our argument involves first showing the existence of a set $\mathcal{X}$ of matrices, so that for each $\boldsymbol{X}^{(i)} \ne \boldsymbol{X}^{(j)} \in \mathcal{X}$, $\|\boldsymbol{X}^{(i)} - \boldsymbol{X}^{(j)}\|_F$ is large. We will imagine obtaining measurements of a randomly chosen matrix in $\mathcal{X}$ and then running an arbitrary recovery procedure. If the recovered matrix is sufficiently close to the original matrix, then we could determine which element of $\mathcal{X}$ was chosen. However, Fano's inequality will imply that the probability of correctly identifying the chosen matrix is small, which will induce a lower bound on how close the recovered matrix can be to the original matrix.

In the proofs of Theorems 3, 4, and 5, we will assume without loss of generality that $d_2 \ge d_1$. Before providing these proofs, however, we first consider the construction of the set $\mathcal{X}$.

4.3.1 Packing set construction

Lemma 3. *Let K be defined as in (11), let $\gamma \le 1$ be such that $\frac{r}{\gamma^2}$ is an integer, and suppose that $\frac{r}{\gamma^2} \le d_1$. There is a set $\mathcal{X} \subset K$ with*

$$|\mathcal{X}| \ge \exp\left(\frac{r d_2}{16\gamma^2}\right)$$

with the following properties:

1. *For all $\boldsymbol{X} \in \mathcal{X}$, each entry has $|X_{i,j}| = \alpha\gamma$.*
2. *For all $\boldsymbol{X}^{(i)}, \boldsymbol{X}^{(j)} \in \mathcal{X}$, $i \ne j$,*

$$\|\boldsymbol{X}^{(i)} - \boldsymbol{X}^{(j)}\|_F^2 > \frac{\alpha^2\gamma^2 d_1 d_2}{2}.$$

Proof. We use a probabilistic argument. The set $\mathcal{X}$ will by constructed by drawing

$$|\mathcal{X}| = \left\lceil \exp\left(\frac{rd_2}{16\gamma^2}\right)\right\rceil \tag{24}$$

matrices $\boldsymbol{X}$ independently from the following distribution. Set $B = \frac{r}{\gamma^2}$. The matrix will consist of blocks of dimensions $B \times d_2$, stacked on top of each other. The entries of the first block (that is, $X_{i,j}$ for $(i,j) \in [B] \times [d_2]$) will be i.i.d. symmetric random variables with values $\pm\alpha\gamma$. Then $\boldsymbol{X}$ will be filled out by copying this block as many times as will fit. That is,

$$X_{i,j} := X_{i',j} \quad \text{where} \quad i' = i \pmod B + 1.$$

Now we argue that with nonzero probability, this set will have all the desired properties. For $\boldsymbol{X} \in \mathcal{X}$,

$$\|\boldsymbol{X}\|_\infty = \alpha\gamma \le \alpha.$$

Further, because $\operatorname{rank} \boldsymbol{X} \le B$,

$$\|\boldsymbol{X}\|_* \le \sqrt{B}\,\|\boldsymbol{X}\|_F = \sqrt{\frac{r}{\gamma^2}}\sqrt{d_1 d_2}\alpha\gamma = \alpha\sqrt{r d_1 d_2}.$$

Thus $\mathcal{X} \subset K$, and all that remains is to show that $\mathcal{X}$ satisfies requirement 2.

For $\boldsymbol{X}, \boldsymbol{W}$ drawn from the above distribution,

$$\begin{aligned}
\|\boldsymbol{X} - \boldsymbol{W}\|_F^2 &= \sum_{i,j}(X_{i,j} - W_{i,j})^2 \\
&\ge \left\lfloor \frac{d_1}{B} \right\rfloor \sum_{i\in[B]} \sum_{j\in[d_2]} (X_{i,j} - W_{i,j})^2 \\
&= 4\alpha^2\gamma^2 \left\lfloor \frac{d_1}{B} \right\rfloor \sum_{i\in[B]} \sum_{j\in[d_2]} \delta_{i,j} \\
&=: 4\alpha^2\gamma^2 \left\lfloor \frac{d_1}{B} \right\rfloor Z(\boldsymbol{X}, \boldsymbol{W}).
\end{aligned}$$

where the $\delta_{i,j}$ are independent 0/1 Bernoulli random variables with mean 1/2. By Hoeffding's inequality and a union bound,

$$\mathbb{P}\left(\min_{\boldsymbol{X}\neq\boldsymbol{W}\in\mathcal{X}} Z(\boldsymbol{X}, \boldsymbol{W}) \le \frac{d_2 B}{4}\right) \le \binom{|\mathcal{X}|}{2} \exp(-Bd_2/8).$$

One can check that for $\boldsymbol{X}$ of the size given in (24), the right-hand side of the above tail bound is less than 1, and thus the event that $Z(\boldsymbol{X}, \boldsymbol{W}) > d_2 B/4$ for all $\boldsymbol{X} \neq \boldsymbol{W} \in \mathcal{X}$ has non-zero probability. In this event,

$$\|\boldsymbol{X} - \boldsymbol{W}\|_F^2 > \alpha^2\gamma^2 \left\lfloor \frac{d_1}{B} \right\rfloor d_2 B \ge \frac{\alpha^2\gamma^2 d_1 d_2}{2},$$

where the second inequality uses the assumption that $d_1 \ge B$ and the fact that $\lfloor x \rfloor \ge x/2$ for all $x \ge 1$. Hence, requirement (2) holds with nonzero probability and thus the desired set exists. □

4.3.2 Proof of Theorem 3

Before we prove Theorem 3, we will need the following lemma about the KL divergence.

Lemma 4. *Suppose that $x, y \in (0,1)$. Then*

$$D(x\|y) \le \frac{(x-y)^2}{y(1-y)}.$$

Proof. Without loss of generality, we may assume that $x < y$. Indeed, $D(1-x\|1-y) = D(x\|y)$, and either $x < y$ or $1-x < 1-y$. Let $z = y - x$. A simple computation shows that

$$\frac{\partial}{\partial z} D(x\|x+z) = \frac{z}{(x+z)(1-x-z)}.$$

Thus, by Taylor's theorem, there is some $\xi \in [0, z]$ so that

$$D(x\|y) = D(x\|x) + z\left(\frac{\xi}{(x+\xi)(1-x-\xi)}\right).$$

Since the right hand side is increasing in ξ, we may replace ξ with z and conclude

$$D(x\|y) \le \frac{(y-x)^2}{y(1-y)},$$

as desired. $\square$

Now, for the proof of Theorem 3, we choose ϵ so that

$$\epsilon^2 = \min\left\{\frac{1}{1024}, C_2\alpha\sqrt{\beta_{3\alpha/4}}\sqrt{\frac{rd_2}{m}}\right\}, \tag{25}$$

where C_2 is an absolute constant to be specified later. We will next use Lemma 3 to construct a set $\mathcal{X}$, choosing γ so that $\frac{r}{\gamma^2}$ is an integer and

$$\frac{4\sqrt{2}\epsilon}{\alpha} \le \gamma \le \frac{8\epsilon}{\alpha}.$$

We can make such a choice because $\epsilon \le 1/32$ and $r \ge 4$. We verify that such a choice for γ satisfies the requirements of Lemma 3. Indeed, since $\epsilon \le 1/32$ and $\alpha \ge 1$ we have $\gamma \le 1/4 < 1$. Further, we assume in the theorem that the right-hand side of (25) is larger than $Cr\alpha^2/d_1$ which implies that $r/\gamma^2 \le d_1$ for an appropriate choice of C.

Let $\mathcal{X}'_{\alpha/2,\gamma}$ be the set whose existence is guaranteed by Lemma 3 with this choice of γ, and with $\alpha/2$ instead of α. We will construct $\mathcal{X}$ by setting

$$\mathcal{X} := \left\{\boldsymbol{X}' + \alpha\left(1 - \frac{\gamma}{2}\right)\mathbf{1} \;:\; \boldsymbol{X}' \in \mathcal{X}'_{\alpha/2,\gamma}\right\}$$

Note that $\mathcal{X}$ has the same size as $\mathcal{X}'_{\alpha/2,\gamma}$, i.e., $|\mathcal{X}|$ satisfies (24). $\mathcal{X}$ also has the same bound on pairwise distances

$$\|\boldsymbol{X}^{(i)} - \boldsymbol{X}^{(j)}\|_F^2 \ge \frac{\alpha^2\gamma^2 d_1 d_2}{8} \ge 4 d_1 d_2 \epsilon^2, \tag{26}$$

and every entry of $\boldsymbol{X} \in \mathcal{X}$ has

$$|X_{i,j}| \in \{\alpha, \alpha'\},$$

where $\alpha' = (1-\gamma)\alpha$. Further, since for $\boldsymbol{X}' \in \mathcal{X}'_{\alpha/2,\gamma}$,

$$\left\|\boldsymbol{X}' + \alpha(1-\gamma/2)\mathbf{1}\right\|_* \le \left\|\boldsymbol{X}'\right\|_* + \alpha(1-\gamma/2)\sqrt{d_1 d_2} \le \alpha\sqrt{r d_1 d_2}$$

for $r \ge 4$ as in the theorem statement.

Now suppose for the sake of a contradiction that there exists an algorithm such that for any $\boldsymbol{X} \in K$, when given access to the measurements $\boldsymbol{Y}_\Omega$, returns an $\widehat{\boldsymbol{X}}$ such that

$$\frac{1}{d_1 d_2}\|\boldsymbol{X} - \widehat{\boldsymbol{X}}\|_F^2 < \epsilon^2 \tag{27}$$

with probability at least 1/4. We will imagine running this algorithm on a matrix $\boldsymbol{X}$ chosen uniformly at random from $\mathcal{X}$. Let

$$\boldsymbol{X}^* = \arg\min_{\boldsymbol{X}^{(i)} \in \mathcal{X}} \|\boldsymbol{X}^{(i)} - \widehat{\boldsymbol{X}}\|_F^2.$$

It is easy to check that if (27) holds, then $\boldsymbol{X}^* = \boldsymbol{X}$. Indeed, for any $\boldsymbol{X}' \in \mathcal{X}$ with $\boldsymbol{X}' \neq \boldsymbol{X}$, from (27) and (26) we have that

$$\|\boldsymbol{X}' - \widehat{\boldsymbol{X}}\|_F = \|\boldsymbol{X}' - \boldsymbol{X} + \boldsymbol{X} - \widehat{\boldsymbol{X}}\|_F \geq \|\boldsymbol{X}' - \boldsymbol{X}\|_F - \|\boldsymbol{X} - \widehat{\boldsymbol{X}}\|_F > 2\sqrt{d_1 d_2}\epsilon - \sqrt{d_1 d_2}\epsilon = \sqrt{d_1 d_2}\epsilon.$$

At the same time, since $\boldsymbol{X} \in \mathcal{X}$ is a candidate for $\boldsymbol{X}^*$, we have that

$$\|\boldsymbol{X}^* - \widehat{\boldsymbol{X}}\|_F \leq \|\boldsymbol{X} - \widehat{\boldsymbol{X}}\|_F \leq \sqrt{d_1 d_2}\epsilon.$$

Thus, if (27) holds, then $\|\boldsymbol{X}^* - \widehat{\boldsymbol{X}}\|_F < \|\boldsymbol{X}' - \widehat{\boldsymbol{X}}\|_F$ for any $\boldsymbol{X}' \in \mathcal{X}$ with $\boldsymbol{X}' \neq \boldsymbol{X}$, and hence we must have $\boldsymbol{X}^* = \boldsymbol{X}$. By assumption, (27) holds with probability at least 1/4, and thus

$$\mathbb{P}\left(\boldsymbol{X} \neq \boldsymbol{X}^*\right) \leq \frac{3}{4}. \tag{28}$$

We will show that this probability must in fact be large, generating our contradiction.

By a variant of Fano's inequality

$$\mathbb{P}\left(\boldsymbol{X} \neq \boldsymbol{X}^*\right) \geq 1 - \frac{\max_{\boldsymbol{X}^{(k)} \neq \boldsymbol{X}^{(\ell)}} D(\boldsymbol{Y}_\Omega | \boldsymbol{X}^{(k)} \parallel \boldsymbol{Y}_\Omega | \boldsymbol{X}^{(\ell)}) + 1}{\log |\mathcal{X}|}. \tag{29}$$

Because each entry of $\boldsymbol{Y}$ is independent,[3]

$$D := D(\boldsymbol{Y}_\Omega | \boldsymbol{X}^{(k)} \parallel \boldsymbol{Y}_\Omega | \boldsymbol{X}^{(\ell)}) = \sum_{(i,j) \in \Omega} D(Y_{i,j} | X_{i,j}^{(k)} \parallel Y_{i,j} | X_{i,j}^{(\ell)}).$$

Each term in the sum is either 0, $D(\alpha \| \alpha')$, or $D(\alpha' \| \alpha)$. By Lemma 4, all of these are bounded above by

$$D(Y_{i,j} | X_{i,j}^{(k)} \parallel Y_{i,j} | X_{i,j}^{(\ell)}) \leq \frac{(f(\alpha) - f(\alpha'))^2}{f(\alpha')(1 - f(\alpha'))},$$

and so, from the intermediate value theorem, for some $\xi \in [\alpha', \alpha]$,

$$D \leq m \frac{(f(\alpha) - f(\alpha'))^2}{f(\alpha')(1 - f(\alpha'))} \leq m \frac{(f'(\xi))^2 (\alpha - \alpha')^2}{f(\alpha')(1 - f(\alpha'))}.$$

Using the assumption that $f'(x)$ is decreasing for $x > 0$ and the definition of $\alpha' = (1 - \gamma)\alpha$, we have

$$D \leq \frac{m(\gamma\alpha)^2}{\beta_{\alpha'}} \leq \frac{64 m \epsilon^2}{\beta_{\alpha'}}.$$

Then from (29) and (28),

$$\frac{1}{4} \leq 1 - \mathbb{P}\left(\boldsymbol{X} \neq \boldsymbol{X}^*\right) \leq \frac{D + 1}{\log |\mathcal{X}|} \leq 16\gamma^2 \left(\frac{\frac{64 m \epsilon^2}{\beta_{\alpha'}} + 1}{r d_2} \right) \leq 1024 \epsilon^2 \left(\frac{\frac{64 m \epsilon^2}{\beta_{\alpha'}} + 1}{\alpha^2 r d_2} \right). \tag{30}$$

We now show that for appropriate values of C_0 and C_2, this leads to a contradiction. First suppose that $64 m \epsilon^2 \leq \beta_{\alpha'}$. In this case we have

$$\frac{1}{4} \leq 1024 \epsilon^2 \frac{2}{\alpha^2 r d_2},$$

[3]Note that here, to be consistent with the literature we are referencing regarding Fano's inequality, D is defined slightly differently than elsewhere in the paper where we would weight D by $1/d_1 d_2$.

which together with (25) implies that $\alpha^2 rd_2 \le 8$. If we set $C_0 > 8$ in (10), then this would lead to a contradiction. Thus, suppose now that $64m\epsilon^2 > \beta_{\alpha'}$. Then (30) simplifies to

$$\frac{1}{4} < \frac{1024 \cdot 128 \cdot m\epsilon^4}{\beta_{\alpha'}\alpha^2 rd_2}.$$

Thus

$$\epsilon^2 > \frac{\alpha\sqrt{\beta_{\alpha'}}}{512\sqrt{2}}\sqrt{\frac{rd_2}{m}}.$$

Note β is increasing as a function of α and $\alpha' \ge 3\alpha/4$ (since $\gamma \le 1/4$). Thus, $\beta_{\alpha'} \ge \beta_{3\alpha/4}$. Setting $C_2 \le 1/512\sqrt{2}$ in (25) now leads to a contradiction, and hence (27) must fail to hold with probability at least 3/4, which proves the theorem.

4.3.3 Proof of Theorem 4

Choose ϵ so that

$$\epsilon^2 = \min\left\{\frac{1}{16}, C_2\alpha\sigma\sqrt{\frac{rd_2}{m}}\right\} \tag{31}$$

for an absolute constant C_2 to be determined later. As in the proof of Theorem 3, we will consider running such an algorithm on a random element in a set $\mathcal{X} \subset K$. For our set $\mathcal{X}$, we will use the set whose existence is guaranteed by Lemma 3. We will set γ so that $\frac{r}{\gamma^2}$ is an integer and

$$\frac{2\sqrt{2}\epsilon}{\alpha} \le \gamma \le \frac{4\epsilon}{\alpha}.$$

This is possible since $\epsilon \le 1/4$ and $r, \alpha \ge 1$. One can check that γ satisfies the assumptions of Lemma 3.

Now suppose that $\boldsymbol{X} \in \mathcal{X}$ is chosen uniformly at random, and let $\boldsymbol{Y} = (\boldsymbol{X} + \boldsymbol{Z})|_\Omega$ as in the statement of the theorem. Let $\widehat{\boldsymbol{X}}$ be any estimate of $\boldsymbol{X}$ obtained from $\boldsymbol{Y}_\Omega$. We begin by bounding the mutual information $I(\boldsymbol{X}; \widehat{\boldsymbol{X}})$ in the following lemma (which is analogous to [12, Equation 9.16]).

Lemma 5.

$$I(\boldsymbol{X}; \widehat{\boldsymbol{X}}) \le \frac{m}{2}\log\left(\sigma^2 + \left(\alpha^2\gamma^2\right)\right).$$

Proof. We begin by noting that

$$I(\boldsymbol{X}_\Omega; \boldsymbol{Y}) = h(\boldsymbol{X}_\Omega + \boldsymbol{Z}_\Omega) - h(\boldsymbol{X}_\Omega + \boldsymbol{Z}_\Omega | \boldsymbol{X}_\Omega) = h(\boldsymbol{X}_\Omega + \boldsymbol{Z}_\Omega) - h(\boldsymbol{Z}_\Omega),$$

where h denotes the differential entropy. Let ξ denote a matrix of i.i.d. ± 1 entries. Then

$$h(\boldsymbol{X}_\Omega \circ \xi + \boldsymbol{Z}_\Omega) = h((\boldsymbol{X}_\Omega + \boldsymbol{Z}_\Omega) \circ \xi) \ge h((\boldsymbol{X}_\Omega + \boldsymbol{Z}_\Omega) \circ \xi \mid \xi) = h(\boldsymbol{X}_\Omega + \boldsymbol{Z}_\Omega),$$

and so, letting $\widetilde{\boldsymbol{X}} = \boldsymbol{X} \circ \xi$,

$$I(\boldsymbol{X}_\Omega; \boldsymbol{Y}) \le h(\widetilde{\boldsymbol{X}}_\Omega + \boldsymbol{Z}_\Omega) - h(\boldsymbol{Z}_\Omega).$$

Treating $\widetilde{\boldsymbol{X}}_\Omega + \boldsymbol{Z}_\Omega$ as a random vector of length m, we compute the covariance matrix as

$$\boldsymbol{\Sigma} := \mathbb{E}\left[\operatorname{vec}(\widetilde{\boldsymbol{X}}_\Omega + \boldsymbol{Z}_\Omega)\operatorname{vec}(\widetilde{\boldsymbol{X}}_\Omega + \boldsymbol{Z}_\Omega)^T\right] = \left(\sigma^2 + (\alpha\gamma)^2\right)\boldsymbol{I}_m.$$

By Theorem 8.6.5 in [12],

$$h(\widetilde{\boldsymbol{X}}_\Omega + \widetilde{\boldsymbol{Z}}_\Omega) \le \frac{1}{2}\log\left((2\pi e)^m \det(\boldsymbol{\Sigma})\right) = \frac{1}{2}\log\left((2\pi e)^m(\sigma^2 + (\alpha\gamma)^2)^m\right).$$

We have that $h(\boldsymbol{Z}_\Omega) = \frac{1}{2}\left((2\pi e)^m \sigma^{2m}\right)$, and so

$$I(\boldsymbol{X}_\Omega; \boldsymbol{Y}) \leq \frac{m}{2} \log\left(1 + \left(\frac{\alpha\gamma}{\sigma}\right)^2\right).$$

Then the data processing inequality implies

$$I(\boldsymbol{X}; \widehat{\boldsymbol{X}}) \leq \frac{m}{2} \log\left(1 + \left(\frac{\alpha\gamma}{\sigma}\right)^2\right),$$

which establishes the lemma. □

We now proceed by using essentially the same argument as in the proof of Theorem 3. Specifically, we suppose for the sake of a contradiction that there exists an algorithm such that for any $\boldsymbol{X} \in K$, when given access to the measurements $\boldsymbol{Y}_\Omega$, returns an $\widehat{\boldsymbol{X}}$ such that

$$\frac{1}{d_1 d_2} \|\boldsymbol{X} - \widehat{\boldsymbol{X}}\|_F^2 < \epsilon^2 \tag{32}$$

with probability at least 1/4. As before, if we set

$$\boldsymbol{X}^* = \arg\min_{\boldsymbol{X}^{(i)} \in \mathcal{X}} \|\boldsymbol{X}^{(i)} - \widehat{\boldsymbol{X}}\|_F^2$$

then we can show that if (32) holds, then $\boldsymbol{X}^* = \boldsymbol{X}$. Thus, if (32) holds with probability at least 1/4 then

$$\mathbb{P}\left(\boldsymbol{X} \neq \boldsymbol{X}^*\right) \leq \frac{3}{4}. \tag{33}$$

However, by Fano's inequality, the probability that $\boldsymbol{X} \neq \widehat{\boldsymbol{X}}$ is at least

$$\mathbb{P}\left(\boldsymbol{X} \neq \widehat{\boldsymbol{X}}\right) \geq \frac{H(\boldsymbol{X}|\widehat{\boldsymbol{X}}) - 1}{\log(|\mathcal{X}|)} = \frac{H(\boldsymbol{X}) - I(\boldsymbol{X}; \widehat{\boldsymbol{X}}) - 1}{\log(|\mathcal{X}|)} \geq 1 - \frac{I(\boldsymbol{X}; \widehat{\boldsymbol{X}}) + 1}{\log|\mathcal{X}|}$$

Plugging in $|\mathcal{X}|$ from Lemma 3 and $I(\boldsymbol{X}; \widehat{\boldsymbol{X}})$ from Lemma 5, and using the inequality $\log(1+z) \leq z$, we obtain

$$\mathbb{P}\left(\boldsymbol{X} \neq \widehat{\boldsymbol{X}}\right) \geq 1 - \frac{16\gamma^2}{r d_2}\left(\frac{m}{2}\left(\frac{\alpha\gamma}{\sigma}\right)^2 + 1\right).$$

Combining this with (33) and using the fact that $\gamma \leq 4\epsilon/\alpha$, we obtain

$$\frac{1}{4} \leq \frac{256\epsilon^2}{\alpha^2 r d_2}\left(8m\left(\frac{\epsilon^2}{\sigma^2}\right) + 1\right).$$

We now argue, as before, that this leads to a contradiction. Specifically, if $8m\epsilon^2/\sigma^2 \leq 1$, then together with (31) this implies that $\alpha^2 r d_2 \leq 128$. If we set $C_0 > 128$ in (10), then this would lead to a contradiction. Thus, suppose now that $8m\epsilon^2/\sigma^2 > 1$, in which case we have

$$\epsilon^2 > \frac{\alpha\sigma}{128}\sqrt{\frac{r d_2}{m}}.$$

Thus, setting $C_2 \leq 1/128$ in (31) leads to a contradiction, and hence (32) must fail to hold with probability at least 3/4, which proves the theorem.

4.3.4 Proof of Theorem 5

The proof of Theorem 5 also mirrors the proof of Theorem 3. The main difference is the observation that the set constructed in Lemma 3 also works with the Hellinger distance. We begin as before by choosing ϵ so that

$$\epsilon^2 = \min\left\{\frac{c}{16}, C_2 \frac{\alpha}{L_1}\sqrt{\frac{rd_2}{m}}\right\}, \tag{34}$$

where C_2 is an absolute constant to be determined. Set γ to be an integer so that

$$2\sqrt{2}\frac{\epsilon}{\alpha c} \leq \gamma \leq \frac{4\epsilon}{\alpha c}.$$

This is possible since by assumption $\alpha \geq 1$ and $\epsilon \leq \frac{c}{4}$. One can check that γ satisfies the assumptions of Lemma 3.

As in the proof of Theorem 3, we will consider running such an algorithm on a random element in a set $\mathcal{X} \subset K$. For our set $\mathcal{X}$, we will use the set whose existence is guaranteed by Lemma 3. Note that since the Hellinger distance is bounded below by the Frobenius norm, we have that for all $\boldsymbol{X}^{(i)} \neq \boldsymbol{X}^{(j)} \in \mathcal{X}$,

$$d_H^2(f(\boldsymbol{X}^{(i)}) - f(\boldsymbol{X}^{(j)})) \geq \|f(\boldsymbol{X}^{(i)}) - f(\boldsymbol{X}^{(j)})\|_F^2 \geq c^2\|\boldsymbol{X}^{(i)} - \boldsymbol{X}^{(j)}\|_F^2 > \frac{c^2}{2}\alpha^2\gamma^2 d_1 d_2 \geq 4 d_1 d_2 \epsilon^2.$$

Now suppose for the sake of a contradiction that there exists an algorithm such that for any $\boldsymbol{X} \in K$, when given access to the measurements $\boldsymbol{Y}_\Omega$, returns an $\widehat{\boldsymbol{X}}$ such that

$$d_H^2(f(\boldsymbol{X}), f(\widehat{\boldsymbol{X}})) < \epsilon^2 \tag{35}$$

with probability at least 1/4. If we set

$$\boldsymbol{X}^* = \arg\min_{\boldsymbol{X}^{(i)} \in \mathcal{X}} d_H^2(f(\boldsymbol{X}^{(i)}) - f(\widehat{\boldsymbol{X}}))$$

then we can show that if (35) holds, then $\boldsymbol{X}^* = \boldsymbol{X}$. Thus, if (35) holds with probability at least 1/4 then

$$\mathbb{P}(\boldsymbol{X} \neq \boldsymbol{X}^*) \leq \frac{3}{4}. \tag{36}$$

However, we may again apply Fano's inequality as in (29). Using Lemma 4 we have

$$D(Y_{i,j}|X_{i,j}^{(k)} \,\|\, Y_{i,j}|X_{i,j}^{(\ell)}) \leq \frac{(f(\alpha\gamma) - f(-\alpha\gamma))^2}{f(\alpha\gamma)(1 - f(\alpha\gamma))} \leq \frac{4(f'(\xi))^2\alpha^2\gamma^2}{f(\alpha\gamma)(1 - f(\alpha\gamma))} \leq \frac{4f^2(\xi)L_{\alpha\gamma}^2\alpha^2\gamma^2}{f(\alpha\gamma)(1 - f(\alpha\gamma))},$$

for some $|\xi| \leq \alpha\gamma$, where $L_{\alpha\gamma}$ is as in (5). By the assumption that $c' < |f(x)| < 1 - c'$ for $|x| < 1$, and that

$$\alpha\gamma \leq \alpha\left(\frac{4\epsilon}{\alpha c}\right) \leq \frac{4\epsilon}{c} \leq 1,$$

we obtain

$$D(Y_{i,j}|X_{i,j}^{(k)} \,\|\, Y_{i,j}|X_{i,j}^{(\ell)}) \leq \frac{4c' L_1^2\alpha^2\gamma^2}{1 - c'} \leq C' L_1^2 \epsilon^2,$$

where $C' = 64c'/(c^2(1 - c'))$. Thus, from (29), we have

$$\frac{1}{4} \leq \frac{C' m L_1^2 \epsilon^2 + 1}{\log|\mathcal{X}|} \leq \frac{256}{c^2}\epsilon^2\left(\frac{C' m L_1^2\epsilon^2 + 1}{\alpha^2 r d_2}\right).$$

We now argue once again that this leads to a contradiction. Specifically, if $C' m L_1^2 \epsilon^2 \leq 1$, then together with (34) this implies that $\alpha^2 r d_2 \leq 128/c$. If we set $C_0 > 128/c$ in (10), then this would lead to a contradiction. Thus, suppose now that $C' m L_1^2 \epsilon^2 > 1$, in which case we have

$$\epsilon^2 > \frac{c}{32\sqrt{2C'}}\frac{\alpha}{L_1}\sqrt{\frac{rd_2}{m}}.$$

Thus setting $C_2 \leq c/32\sqrt{2C'}$ in (34) leads to a contradiction, and hence (35) must fail to hold with probability at least 3/4, which proves the theorem.

5 Discussion

Many of the applications of matrix completion consider discrete data, sometimes consisting of binary measurements. This paper addresses such situations. However, matrix completion from noiseless binary measurements is extremely ill-posed, even if one collects a binary measurement from all of the matrix entries. Fortunately, when there are some stochastic variations (noise) in the problem, matrix reconstruction becomes well-posed. We demonstrate that the unknown matrix can be accurately and efficiently recovered from binary measurements. When the infinity norm of the unknown matrix is bounded by a constant, we show that our error bounds are tight to within a constant and even match what is possible for undiscretized data. We also show that the binary probability distribution can be reconstructed over the entire matrix without any assumption on the infinity-norm, and we give a matching lower bound (up to a constant).

Our theory considers approximately low-rank matrices—in particular, we assume that the singular values belong to a scaled Schatten-1 ball. It would be interesting to see whether more accurate reconstruction could be achieved under the assumption that the unknown matrix has precisely r nonzero singular values. It would also be interesting to study whether our ideas can be extended to deal with measurements that are quantized to more than 2 (but still a small number) of different values.

References

[1] A. Ai, A. Lapanowski, Y. Plan, and R. Vershynin. One-bit compressed sensing with non-gaussian measurements. *Arxiv preprint arxiv:1208.6279*, 2012.

[2] F. Bach. Self-concordant analysis for logistic regression. *Elec. J. Stat.*, 4:384–414, 2010.

[3] P. Biswas, T.-C. Lian, T.-C. Wang, and Y. Ye. Semidefinite programming based algorithms for sensor network localization. *ACM Trans. Sen. Netw.*, 2(2):188–220, 2006.

[4] I. Borg and P. Groenen. *Modern multidimensional scaling*. Springer Science, New York, NY, 2010.

[5] P. Boufounos and R. Baraniuk. 1-Bit compressive sensing. In *Proc. IEEE Conf. Inform. Science and Systems (CISS)*, Princeton, NJ, Mar. 2008.

[6] F. Bunea. Honest variable selection in linear and logistic regression models via ℓ_1 and $\ell_1 + \ell_2$ penalization. *Elec. J. Stat.*, 2:1153–1194, 2008.

[7] E. Candès. Compressive sampling. In *Proc. Int. Congress of Math.*, Madrid, Spain, Aug. 2006.

[8] E. Candès and M. Davenport. How well can we estimate a sparse vector? *to appear in Appl. Comput. Harmon. Anal.*, 2012.

[9] E. Candès and Y. Plan. Matrix completion with noise. *Proc. IEEE*, 98(6):925–936, 2010.

[10] E. Candès and B. Recht. Exact matrix completion via convex optimization. *Found. Comput. Math.*, 9(6):717–772, 2009.

[11] E. Candès and T. Tao. The power of convex relaxation: Near-optimal matrix completion. *IEEE Trans. Inform. Theory*, 56(5):2053–2080, 2010.

[12] T. Cover and J. Thomas. *Elements of information theory*. Wiley-Interscience, New York, NY, 1991.

[13] M. Davenport, M. Duarte, Y. Eldar, and G. Kutyniok. Introduction to compressed sensing. In Y. Eldar and G. Kutyniok, editors, *Compressed Sensing: Theory and Applications*. Cambridge University Press, Cambridge, UK, 2012.

[14] D. Donoho. Compressed sensing. *IEEE Trans. Inform. Theory*, 52(4):1289–1306, 2006.

[15] S. Gaïffas and G. Lecué. Sharp oracle inequalities for the prediction of a high-dimensional matrix. *Arxiv preprint arxiv:1008.4886*, 2010.

[16] D. Gleich and L.-H. Lim. Rank aggregation via nuclear norm minimization. In *Proc. ACM SIGKDD Int. Conf. on Knowledge, Discovery, and Data Mining (KDD)*, San Diego, CA, Aug. 2011.

[17] D. Goldberg, D. Nichols, B. M. Oki, and D. Terry. Using collaborative filtering to weave an information tapestry. *Comm. ACM*, 35(12):61–70, 1992.

[18] P. Green and Y. Wind. *Multivariate decisions in marketing: A measurement approach*. Dryden, Hinsdale, IL, 1973.

[19] D. Gross. Recovering low-rank matrices from few coefficients in any basis. *IEEE Trans. Inform. Theory*, 57(3):1548–1566, 2011.

[20] D. Gross, Y. Liu, S. Flammia, S. Becker, and J. Eisert. Quantum state tomography via compressed sensing. *Physical Review Letters*, 105(15):150401, 2010.

[21] A. Gupta, R. Nowak, and B. Recht. Sample complexity for 1-bit compressed sensing and sparse classification. In *Proc. IEEE Int. Symp. Inform. Theory (ISIT)*, Austin, TX, June 2010.

[22] L. Jacques, J. Laska, P. Boufounos, and R. Baraniuk. Robust 1-bit compressive sensing via binary stable embeddings of sparse vectors. *Arxiv preprint arxiv:1104.3160*, 2011.

[23] S. Kakade, O. Shamir, K. Sridharan, and A. Tewari. Learning exponential families in high-dimensions: Strong convexity and sparsity. In *Proc. Int. Conf. Art. Intell. Stat. (AISTATS)*, Clearwater Beach, FL, Apr. 2009.

[24] R. Keshavan, A. Montanari, and S. Oh. Matrix completion from a few entries. *IEEE Trans. Inform. Theory*, 56(6):2980–2998, 2010.

[25] R. Keshavan, A. Montanari, and S. Oh. Matrix completion from noisy entries. *J. Machine Learning Research*, 11:2057–2078, 2010.

[26] O. Klopp. High dimensional matrix estimation with unknown variance of the noise. *Arxiv preprint arxiv:1112.3055*, 2011.

[27] O. Klopp. Rank penalized estimators for high-dimensional matrices. *Elec. J. Stat.*, 5:1161–1183, 2011.

[28] V. Koltchinskii. Von Neumann entropy penalization and low-rank matrix estimation. *Ann. Stat.*, 39(6):2936–2973, 2012.

[29] V. Koltchinskii, K. Lounici, and A. Tsybakov. Nuclear-norm penalization and optimal rates for noisy low-rank matrix completion. *Ann. Stat.*, 39(5):2302–2329, 2011.

[30] J. Laska and R. Baraniuk. Regime change: Bit-depth versus measurement-rate in compressive sensing. *IEEE Trans. Signal Processing*, 60(7):3496–3505, 2012.

[31] M. Ledoux and M. Talagrand. *Probability in Banach Spaces: Isoperimetry and Processes*. Springer-Verlag, Berlin, Germany, 1991.

[32] Z. Liu and L. Vandenberghe. Interior-point method for nuclear norm approximation with application to system identification. *SIAM J. Matrix Analysis and Applications*, 31(3):1235–1256, 2009.

[33] L. Meier, S. Van De Geer, and P. Bühlmann. The group lasso for logistic regression. *J. Royal Stat. Soc. B*, 70(1):53–71, 2008.

[34] G. Miller. The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Rev.*, 63(2):81–97, 1956.

[35] S. Negahban, P. Ravikumar, M. Wainwright, and B. Yu. A unified framework for high-dimensional analysis of M-estimators with decomposable regularizers. *Arxiv preprint arxiv:1010.2731*, 2010.

[36] S. Negahban and M. Wainwright. Restricted strong convexity and weighted matrix completion: Optimal bounds with noise. *J. Machine Learning Research*, 13:1665–1697, 2012.

[37] Y. Plan and R. Vershynin. One-bit compressed sensing by linear programming. *Arxiv preprint arxiv:1109.4299*, 2011.

[38] Y. Plan and R. Vershynin. Robust 1-bit compressed sensing and sparse logistic regression: A convex programming approach. *Arxiv preprint arxiv:1202.1212*, 2012.

[39] G. Raskutti, M. Wainwright, and B. Yu. Minimax rates of estimation for high-dimensional linear regression over ℓ_q-balls. *IEEE Trans. Inform. Theory*, 57(10):6976–6994, 2011.

[40] P. Ravikumar, M. Wainwright, and J. Lafferty. High-dimensional Ising model selection using ℓ_1-regularized logistic regression. *Ann. Stat.*, 38(3):1287–1319, 2010.

[41] B. Recht. A simpler approach to matrix completion. *J. Machine Learning Research*, 12:3413–3430, 2011.

[42] A. Rohde and A. Tsybakov. Estimation of high-dimensional low-rank matrices. *Ann. Stat.*, 39(2):887–930, 2011.

[43] Y. Seginer. The expected norm of random matrices. *Combinatorics, Probability, and Computing*, 9(2):149–166, 2000.

[44] A. Singer. A remark on global positioning from local distances. *Proc. Natl. Acad. Sci.*, 105(28):9507–9511, 2008.

[45] A. Singer and M. Cucuringu. Uniqueness of low-rank matrix completion by rigidity theory. *SIAM J. Matrix Analysis and Applications*, 31(4):1621–1641, 2010.

[46] I. Spence and D. Domoney. Single subject incomplete designs for nonmetric multidimensional scaling. *Psychometrika*, 39(4):469–490, 1974.

[47] S. Van De Geer. High-dimensional generalized linear models and the lasso. *Ann. Stat.*, 36(2):614–645, 2008.

[48] A. Zymnis, S. Boyd, and E. Candès. Compressed sensing with quantized measurements. *IEEE Signal Processing Lett.*, 17(2):149–152, 2010.