Robust Matrix Completion from Quantized Observations

Jie Shen
Stevens Institute of Technology
New Jersey, USA

Pranjal Awasthi
Rutgers University
New Jersey, USA

Ping Li
Baidu Research
Washington, USA

Abstract

1-bit matrix completion refers to the problem of recovering a real-valued low-rank matrix from a small fraction of its sign patterns. In many real-world applications, however, the observations are not only highly quantized, but also grossly corrupted. In this work, we consider the noisy statistical model where each observed entry can be flipped with some probability after quantization. We propose a simple maximum likelihood estimator which is shown to be robust to the sign flipping noise. In particular, we prove an upper bound on the statistical error, showing that with overwhelming probability $n = \mathcal{O}\left(\mathrm{poly}(1 - 2\,\mathbb{E}[\tau])^{-2} rd \log d\right)$ samples are sufficient for accurate recovery, where r and d are the rank and dimension of the underlying matrix respectively, and $\tau \in [0, 1/2)$ is a random variable that parameterizes the sign flipping noise. Furthermore, a lower bound is established showing that the obtained sample complexity is near-optimal for prevalent statistical models. Finally, we substantiate our theoretical findings with a comprehensive study on synthetic and realistic data sets, and demonstrate the state-of-the-art performance.

1 Introduction

Many practical problems can be formulated as recovering an incomplete matrix from the small portion of its components, which is known as matrix completion [10]. For instance, in the *Netflix Prize* competition, the underlying matrix consists of movie ratings from a variety of users, and the task is to predict the taste of the users for their unrated movies. This problem has been studied for a decade, and the matrix factorization framework was proposed as an early answer [46]. In the seminal work [10], it was shown that if the singular vectors of the matrix to be recovered are dense enough and the observed entries are sampled uniformly random, then with high probability, a simple convex program guarantees exact recovery of the true matrix.

Inspired by the elegant work of [10], a plethora of theoretical results have been established that study the matrix completion problem from different aspects. A partial list includes: understanding and improving the sample complexity [12, 26, 21, 15], addressing structured noise such as outliers [9, 27, 29, 16], developing fast and provable optimization algorithms [24, 25], mitigating practical issues such as memory cost [44, 4, 52], to name just a few. Orthogonal to these work where the observed entries are real-valued, [17] considered the problem in the 1-bit setting. That is, given a target low-rank matrix which is real-valued, one only observes some of its sign patterns ($+1$ or -1) determined by the true matrix. The goal, however, is still to recover the real-valued matrix by using a few samples.

The 1-bit setting is of broad interest for the machine learning community. On one hand, it immediately eases the data acquisition process since it is always a simpler task to ask a user whether he likes a movie or not than having him submit a one to five stars rating. In fact, Netflix recently changed its rating system that only requires the user to say a “thumbs up” or “thumbs down”. From the theoretical perspective, however, it raises the challenge that a straightforward observation model makes the problem ill-posed. To be more detailed, suppose that the binary patterns are obtained by taking the sign of the entries of the true matrix. Then even for a rank-one matrix $M = \boldsymbol{u}\boldsymbol{v}^\top$ where $\boldsymbol{u}$ and $\boldsymbol{v}$ are column vectors, one can freely modify the magnitude of the elements of $\boldsymbol{u}$ and $\boldsymbol{v}$ without changing the sign patterns of M. The second issue coming up with the 1-bit setting is a tractable recovery paradigm. Since the sign function is not convex, one cannot tailor the nuclear-norm based convex pro-

Proceedings of the 22nd International Conference on Artificial Intelligence and Statistics (AISTATS) 2019, Naha, Okinawa, Japan. PMLR: Volume 89.

gram [10] to this case. Another concern is the loss of estimation accuracy owing to quantization. Thus, a precise characterization of the tradeoff between the bit depth and the sample size is crucial. Related to the sign patterns, it is also important to ask if there is a provable algorithm that is tolerant to noise.

[17] gave a partial answer to these questions by showing that, a nuclear-norm based program guarantees accurate recovery from binary measurements if the observations are generated from a distribution parameterized by the true matrix. [5] further derived the statistical error rate of multi-bit quantization. We in this paper tackle the problem of *robust 1-bit matrix completion*: **(a)** can we accurately recover the matrix in polynomial time if the observations are flipped with some probability close to $1/2$; **(b)** if yes, how many samples suffice and is this sample complexity optimal.

Our motivation is two-fold. For practitioners, realistic data are usually discrete. For instance, the data matrix of the social network that represents whether two individuals are friends or not is binary. Sometimes the data are intended to be quantized, due to memory or communication limitation. The system side, nevertheless, has to process this quantized feedback/signal and predict a real value, for example, a likelihood value that the user will watch a new movie. On the other hand, there is a large body of work studying the robustness of standard matrix completion while little is known for the 1-bit case. Unlike the standard problem, the sign flipping noise is no longer additive which poses specific challenges for theoretical analysis.

Contribution. We offer an affirmative answer to the noisy 1-bit matrix completion problem. In particular, we consider the following noise model: for each binary observation, it is flipped with probability $\tau \in [0, 1/2)$ where τ itself is a random variable. This means each entry is flipped with different probability conditioning on τ. We propose a novel nuclear-norm constrained convex program and prove that for any rank-r matrix $M \in \mathbb{R}^{d_1 \times d_2}$ that satisfies certain conditions, it accurately recovers M with high probability, in the sense that the estimation error vanishes when the sample size scales as $\mathcal{O}\left(\mathrm{poly}\,(1 - 2\,\mathbb{E}[\tau])^{-2}\, r(d_1 + d_2)\log(d_1 d_2)\right)$. We also establish a lower bound on the statistical error, showing that the sample complexity is near-optimal.

In our construction of the estimator and the subsequent analysis, we only assume that we have the knowledge of the expected value of τ. This is more practical than requiring the knowledge of its distribution. Interestingly, in our experiments we discover that such assumption can further be relaxed. To be more detailed, we show on realistic data that an estimator constructed with an upper bound of the value of τ performs as well as the one with $\mathbb{E}[\tau]$, suggesting possible extension of our noise model.

1.1 Related Work

Matrix completion is closely related to compressed sensing (CS) [18, 49, 51] where the goal is to recover a sparse vector from its compressed measurements. It is now well-understood that if the sensing matrix satisfies the restricted isometry property [11], then either convex programs such as basis pursuit [14] and Lasso [47, 50] or greedy algorithms like orthogonal matching pursuit [35, 48] or iterative hard thresholding [6, 41, 42, 43] can be used for sparse recovery. Encouraged by the success of compressed sensing, a large body of work was devoted to the nuclear-norm based convex program for low-rank matrix recovery, in view of the analogy between the ℓ_1 norm and the nuclear norm [19, 39, 13]. However, the essential difference is that the sampling operator in compressed sensing is Gaussian, while for matrix completion it is a zero-one matrix $\boldsymbol{e}_i \boldsymbol{e}_j^\top$, where $\boldsymbol{e}_i$ is the ith canonical basis and likewise for $\boldsymbol{e}_j$. In this light, theoretical results in CS cannot be transferred to the matrix completion problem directly [38].

In compressed sensing, the 1-bit setting has received a broad attention due to [7]. There is a variety of appealing work contributed to this emerging field [22, 23, 20], while recently [37] gave an optimal sample complexity that ensures accurate recovery of the direction of the signal. It is very interesting to contrast such a result to the matrix completion problem, where we recall that in the matrix case, even the direction (i.e., $\boldsymbol{u}\boldsymbol{v}^\top/(\|\boldsymbol{u}\|_2 \cdot \|\boldsymbol{v}\|_2)$) cannot be recovered from the knowledge of $\mathrm{sign}\left(\boldsymbol{u}\boldsymbol{v}^\top\right)$. This again suggests discrepancy between compressed sensing and matrix completion. Very recently, the statistical tradeoff between the sample size and bit depth of compressed sensing was investigated in [45], and a guaranteed estimator of the magnitude of the signal was proposed in [28]. The tradeoff of quantized matrix completion was also tackled in [5], but a full picture is still missing.

Of specific interest to the 1-bit setting is the sign flipping noise. Such kind of noise has been widely studied in the learning theory community for more than a decade [32, 3, 53], in the context of learning halfspaces. However, the target vector therein is a general object, i.e., without the sparsity structure. A unified analysis was presented recently in [2], showing possible improvement on noisy 1-bit compressed sensing by using tools from learning theory.

Despite these promising results in 1-bit compressed sensing and learning theory, it turns out that the robustness of 1-bit matrix completion is not well-understood until now. Though these two problems are inherently linked, it has been recognized that extra efforts have to be made in the matrix case. In this work, we take a step to study the model where the noise has the same distribution over the observed entries. This is a popular noise model that was also considered in [37] in the context of 1-bit compressed sensing. Our empirical study, however, sheds light on the more challenging bounded noise model [32].

Notation. We collect the pieces of notation that will be used in this paper. For a matrix $M \in \mathbb{R}^{d_1 \times d_2}$, we use M_{ij} to denote its (i,j)-th entry. The transpose of M is denoted by $M^\top$. There are three matrix norms that will be involved: the Frobenius norm $\|M\|_F := (\sum_{i,j} M_{ij}^2)^{1/2}$, the infinity norm $\|M\|_\infty := \max_{i,j} |M_{ij}|$ where $|x|$ denotes the absolute value of x, and the nuclear norm $\|M\|_* := \sum_{i=1}^r \sigma_i(M)$ where $\sigma_i(M)$ is the ith singular value and r is the rank of M.

Suppose d_1 and d_2 are two positive integers. We write $[d_1] \times [d_2]$ for the index set $\{(i,j) : 1 \le i \le d_1, 1 \le j \le d_2\}$. For a finite set Ω, we slightly abuse the notation to denote its cardinality by $|\Omega|$.

Throughout the paper, f and g are reserved for particular functions. Hence, f' and g' should be interpreted as the derivative evaluated at some point. We also reserve the upright letter C and its subscript variants (e.g., C_0, C_1) for absolute constants, whose values may change from appearance to appearance.

Finally, the sign function $\operatorname{sign}(x)$ outputs $+1$ if $x \ge 0$ and outputs -1 otherwise. For a matrix M, $\operatorname{sign}(M)$ operates in an entry-wise manner. The indicator function is denoted by $\mathbf{1}_{\{E\}}$, which equals one if the event E is true and zero otherwise.

2 Problem Setup

In this section, we formulate the problem. Recall that $M \in \mathbb{R}^{d_1 \times d_2}$ is the underlying low-rank matrix that we aim to recover, and $\Omega \subset [d_1] \times [d_2]$ is a subset that indexing the observed components. In conventional matrix completion [10], one observes M_{ij} for $(i,j) \in \Omega$. Though it seems natural to consider the 1-bit matrix completion problem as a recovery from $\operatorname{sign}(M_\Omega)$, Davenport et al. pointed out that it is not possible even when the matrix M has rank one [17]. The good news is that if we add noise (e.g., Gaussian, logistic) before quantization, it is tractable to solve the problem. Formally, the observation model considered in [17] is as follows: for all $(i,j) \in \Omega$, we observe

$$Y_{ij} = \operatorname{sign}(M_{ij} + Z_{ij}), \tag{1}$$

where $\{Z_{ij}\}$ are i.i.d. random noise. With a proper choice of a differentiable function $f : \mathbb{R} \to [0,1]$, (1) is equivalent to the following probabilistic model:

$$Y_{ij} = \begin{cases} +1, & \text{with probability } f(M_{ij}), \\ -1, & \text{with probability } 1 - f(M_{ij}). \end{cases} \tag{2}$$

In fact, we can set $f(x) = \Pr(Z_{11} + x \ge 0)$ for which the model (1) reduces to the model (2). Conversely, given the function $f(x)$, we may think of $\{Z_{ij}\}$ as i.i.d. random noise with cumulative distribution function $F(x) := \Pr(Z_{11} < x) = 1 - f(-x)$. In this way, (2) reduces to (1).

In this paper, we will mainly consider the model (2), which is viewed as a noiseless probabilistic model. With this in mind, we are in the position to introduce the *noisy* probabilistic model. Our central interest falls into the random sign flipping. That is, in place of observing Y_{ij} as in (2), we observe

$$Y'_{ij} = \delta_{ij} Y_{ij}, \ \forall\ (i,j) \in \Omega, \tag{3}$$

where $\{\delta_{ij}\}_{(i,j)\in\Omega}$ are independent random variables such that given τ

$$\delta_{ij} = \begin{cases} +1, & \text{with probability } 1 - \tau, \\ -1, & \text{with probability } \tau. \end{cases} \tag{4}$$

Above, τ itself might be an unknown random variable but we impose $0 \le \tau < 1/2$ to prevent model ambiguity. Note that $\tau = 0$ corresponds to the noiseless model studied in [17, 5]. The model (3) together with (4) indicate that for each element in Ω, with probability τ the sign is flipped. A remarkable difference between our model and those considered in [37, 45] is that they treat τ as a fixed known parameter which is not realistic in real-world problems.

It is worth mentioning that a more general noise model is that each δ_{ij} is parameterized by τ_{ij}, where $\{\tau_{ij}\}$ may differ from each other but subject to the constraint $0 \le \tau_{ij} \le \tau < 1/2$ for some parameter τ. This is known as bounded noise (or Massart noise) [32] that has received a broad attention in learning theory [1, 3]. We will show in the experiments that the proposed estimator performs well in this situation, and a thorough theoretical study is left as a future work.

Before presenting our estimator for M, we need a few assumptions that were also made in [36, 31, 34, 38].

- $(A1)$ Given $n > 0$, each component (i,j) is included in Ω with probability $\frac{n}{d_1 d_2}$. Hence, $\mathbb{E}|\Omega| = n$.
- $(A2)$ The maximum absolute value of M is upper bounded by a parameter α, i.e., $\|M\|_\infty \le \alpha$.

$(A3)$ M lies in a nuclear-norm ball with radius $\alpha\sqrt{rd_1d_2}$ where r is the rank of M.

Note that $(A1)$ assumes a Bernoulli sampling scheme for Ω which is more convenient to analyze than the uniform sampling. In fact, the equivalence between these two sampling models was pointed out in [12, 8]. The second assumption essentially excludes the case that M is too spiky. Otherwise, the recovery of M is ill-posed [33, 38]. The parameter α is pre-defined, and in practical applications it has to be tuned. Finally, $(A3)$ acts as a convex relaxation to the exact rank constraint $\operatorname{rank}(M) \le r$. To see this, we note that by algebra

$$\|M\|_* \le \sqrt{r}\,\|M\|_F \le \alpha\sqrt{rd_1d_2}.$$

As we will illustrate later, $(A3)$ also allows us to approximate M by solving a *convex* program.

Under these assumptions, we propose to solve the following program in order to approximate M:

$$\begin{aligned}&\max_X\ L_{\Omega,Y'}(X),\\ &\text{s.t. } \|X\|_\infty \le \alpha,\ \|X\|_* \le \alpha\sqrt{rd_1d_2}.\end{aligned} \tag{5}$$

Above, the objective function $L_{\Omega,Y'}(X)$ is given by

$$\begin{aligned}L_{\Omega,Y'}(X) = \sum_{(i,j)\in\Omega} \Big[&\mathbf{1}_{\{Y'_{ij}=1\}}\log g(X_{ij})\\ &+ \mathbf{1}_{\{Y'_{ij}=-1\}}\log\left(1-g(X_{ij})\right)\Big],\end{aligned} \tag{6}$$

where $g(x)$ is the function (to be clarified) such that for every $(i,j)\in\Omega$, Y'_{ij} equals 1 with probability $g(M_{ij})$. In this light, it is not hard to see that $L_{\Omega,Y'}(X)$ is the log-likelihood function and (5) is the maximum likelihood estimator (MLE). We remark that the two constraints in (5) are enforced to accommodate our assumptions $(A2)$ and $(A3)$.

It remains to characterize the function $g(x)$ which is a crucial component of (5). Note that in view of (3) and (4), we have the following conditional probability:

$$\Pr\left(Y'_{ij}=1 \mid \tau\right) = (1-\tau)f(M_{ij}) + \tau(1-f(M_{ij})). \tag{7}$$

Thus, once some statistics of τ is known we are able to evaluate $g(x)$.

(I) τ is discrete. In this case, let us suppose that the random variable τ takes value in $\{\tau_1,\tau_2,\ldots,\tau_s\}$ with corresponding probability $\{p_1,p_2,\ldots,p_s\}$. It then follows that

$$\begin{aligned}\Pr\left(Y'_{ij}=1\right) &= \sum_{k=1}^{s}\Pr\left(Y'_{ij}=1,\tau=\tau_k\right)\\ &= \sum_{k=1}^{s} p_k\Pr\left(Y'_{ij}=1\mid\tau=\tau_k\right).\end{aligned}$$

Hence, letting

$$\begin{aligned}g(x) &= \sum_{k=1}^{s} p_k\left((1-\tau_k)f(x)+\tau_k(1-f(x))\right)\\ &= f(x)\,\mathbb{E}[1-2\tau]+\mathbb{E}[\tau].\end{aligned} \tag{8}$$

gives $\Pr(Y'_{ij}=1) = g(M_{ij})$ as desired.

(II) τ is continuous. Suppose that the probability density function of τ is $h_\tau(\cdot)$. Then by simple calculation, it can be shown that

$$\begin{aligned}g(x) &= \int_t h_\tau(t)\Big[(1-t)f(x)+t(1-f(x))\Big]dt\\ &= f(x)\,\mathbb{E}[1-2\tau]+\mathbb{E}[\tau],\end{aligned} \tag{9}$$

which is identical to the discrete case. Therefore, it turns out that the random flipping noise (4) affects the recovery only through the mean.

3 Main Results

Our main results characterize the statistical rate of the optimum to (5). There are two important quantities that we will need in the theoretical analysis, i.e.,

$$\rho_\gamma^+ \overset{\text{def}}{=} \sup_{|x|\le\gamma}\frac{|g'(x)|}{g(x)(1-g(x))}, \tag{10}$$

$$\rho_\gamma^- \overset{\text{def}}{=} \sup_{|x|\le\gamma}\frac{g(x)(1-g(x))}{(g'(x))^2}. \tag{11}$$

In the above expressions, γ is a generic positive parameter which will be specified to different values in the sequel. It is not hard to see that the quantity ρ_γ^+ is essentially the Lipschitz constant of the log-likelihood function $L_{\Omega,Y'}(X)$. While the quantity ρ_γ^- is not associated with the curvature explicitly, there is still some intuitive explanation on why it enters our analysis. Indeed, presume that $g(x)$ is bounded from below in the interval $[-\gamma,\gamma]$. As $g'(x)$ approaches zero, we find that ρ_γ^- tends to infinity since

$$\frac{\mathrm{C}}{(g'(x))^2} \le \frac{g(x)(1-g(x))}{(g'(x))^2} \le \frac{1}{2(g'(x))^2}$$

for some constant C. In view of (9), this in turn suggests that either the function $f(x)$ is quite flat in the interval or $\mathbb{E}[\tau]$ is close to $1/2$, making it difficult to distinguish the entries of M.

3.1 Upper Bound

With these notions on hand, we state our first result which upper bounds the estimation error of the solution of (5) for the recovery of M.

Theorem 1 (Upper Bound)**.** *Assume (A1), (A2) and (A3). Suppose that the observation model follows* (3)*. Denote $\widehat{M}$ the optimum of* (5)*. Then, with probability at least $1 - \mathrm{C}_1/(d_1+d_2)$, we have*

$$\frac{1}{d_1 d_2}\left\|\widehat{M} - M\right\|_F^2 \leq \psi_\alpha \sqrt{\frac{r(d_1+d_2)}{n}},$$

provided that $n \geq (d_1+d_2)\log(d_1 d_2)$. Above, $\psi_\alpha = \mathrm{C}_2 \rho_\alpha^+ \rho_\alpha^-$.

The theorem implies that as soon as we randomly sample $n \geq \psi_\alpha^2 r(d_1+d_2)\log(d_1 d_2)$ entries, the estimation error vanishes. Note that the dependence on the matrix rank r and the dimension (d_1, d_2) is optimal up to a logarithmic factor. The theorem also suggests that the random flipping noise τ affects the recovery through the quantity ψ_α, which is multiplicative. For concreteness, we give estimates of the quantity ψ_α for prevalent statistical models. In the following we write $a := \mathbb{E}[\tau]$ for brevity.

- Logistic regression: $f(x) = e^x/(1+e^x)$. We have

  $$\rho_\alpha^+ = 1,$$
  $$\rho_\alpha^- = \frac{(1+e^\alpha)^2}{(1-2a)^2 e^{2\alpha}}\left((1-e^\alpha)a + e^\alpha\right)\left((e^\alpha - 1)a + 1\right).$$

  Therefore, if we treat the parameter α as a constant, say $e^\alpha = 2$, it follows that

  $$\psi_\alpha = \mathrm{C}_0 \frac{(a+1)(2-a)}{(1-2a)^2} = \frac{\mathrm{C}_0}{(1-2a)^2},$$

  where the second equality follows by investigating the asymptotic behavior when a approaches $1/2$ from below. The above quickly implies that the sample size $n = \mathcal{O}\left((1-2a)^{-4} r(d_1+d_2)\log(d_1 d_2)\right)$ suffices for accurate recovery even when nearly half of the entries are flipped.

- Probit regression: $f(x) = \Phi(x/\sigma)$. That is, $\{Z_{ij}\}$ in (1) are Gaussian random variables with mean zero and variance σ^2. We have

  $$\rho_\alpha^+ \leq \frac{4}{(1-2a)\sigma}\left(\frac{\alpha}{\sigma} + 1\right),$$
  $$\rho_\alpha^- \leq \frac{\pi\sigma^2}{(1-2a)^2}\exp(\alpha^2/(2\sigma^2)),$$
  $$\psi_\alpha = \mathcal{O}\left(\frac{\alpha+\sigma}{(1-2a)^3}\exp\left(\frac{\alpha^2}{2\sigma^2}\right)\right).$$

  It is not hard to see that there exists a threshold $\sigma^* > \alpha$ that minimizes the right-hand side above, hence is a heuristically optimal choice. When $\sigma < \sigma^*$, one can increase the variance to obtain a better error bound. This is not surprising since on one spectrum, if the variance is too small, the model (1) reduces to $Y_{ij} = \mathrm{sign}\left(M_{ij}\right)$ for which recovery is not possible [17]. On the other extreme, if σ is too large, then the function $f'(x)$ (and hence $g'(x)$) becomes flat, which makes recovery challenging.

  In the regime where the parameter α is a constant, we obtain the sample complexity $n = \mathcal{O}\left((1-2a)^{-6} r(d_1+d_2)\log(d_1 d_2)\right)$ which is worse than the logistic case. These two models are quite similar to each other but Gaussian distribution has a lighter tail, which might be the reason. We leave a concrete study as future work.

- Laplacian: $f'(x) = -\frac{1}{2b}\exp(-\left|x\right|/b)$. We have

  $$\rho_\alpha^+ = \frac{2(1-2a)}{b},$$
  $$\rho_\alpha^- = \frac{b^2}{(1-2a)^2}\left(2(e^{\alpha/b}-1)a + 1\right) \times \left(2(1-e^{\alpha/b})a + 2e^{\alpha/b} - 1\right),$$

  which yields

  $$\psi_\alpha = \mathcal{O}\left(\frac{1}{1-2a}\right),$$

  provided that the parameters α and b are constants, e.g., $e^{\alpha/b} = 2$. This gives us the sample complexity $n = \mathcal{O}\left((1-2a)^{-2} r(d_1+d_2)\log(d_1 d_2)\right)$ which is better than the logistic case. It is also worth mentioning that like probit regression, there exists an optimal choice of the parameter b that minimizes the upper bound of the statistical error, though we do not pursue it here.

3.2 Lower Bound

Our second theorem provides a lower bound on the statistical error for the recovery of M. It asserts that under the observation model (3) and sampling scheme (A1), we can always find an instance M satisfying (A2) and (A3), such that with a non-trivial probability (say, 3/4), any algorithm has to access as many samples as Theorem 1 suggests for recovery.

Theorem 2 (Lower Bound)**.** *Fix the parameters α, r, d_1 and d_2 with $\alpha \geq 1$ and $r \geq 16$. Suppose that $\alpha^2 r \max\{d_1, d_2\} \geq \mathrm{C}_0$ for some absolute constant C_0, and $g'(x)$ is non-increasing for $x > 0$. Let Ω by an arbitrary index set with $|\Omega| = n$ and assume the noisy observation model* (3)*. Then there exists M satisfying (A2) and (A3) such that for any algorithm, with probability at least 3/4, its output $\widehat{M}$ satisfies*

$$\frac{1}{d_1 d_2}\left\|\widehat{M} - M\right\|_F^2 \geq \min\left\{\mathrm{C}_1, \mathrm{C}_2 \phi_\alpha \sqrt{\frac{r\max\{d_1, d_2\}}{n}}\right\},$$

provided that the right-hand side is larger than $r\alpha^2/\min\{d_1, d_2\}$. *Above,* $\phi_\alpha = \alpha(\rho^-_{0.75\alpha})^{1/2}$.

A few remarks are in order. First and foremost, it is shown that $n = \mathcal{O}\left(\phi_\alpha^2 r d_2\right)$ samples are necessary for accurate recovery where we assume $d_1 \leq d_2$ without loss of generality. The dependence on the rank r and matrix dimension (d_1, d_2) matches the upper bound in Theorem 1 (up to a logarithmic factor), justifying the optimality provided that α is a constant. Regarding the noise parameter $\mathbb{E}[\tau]$ contained in ϕ_α, it is not hard to see that for all choices of $f(x)$ (i.e., logistic, probit, laplacian), our lower bound implies that n is proportional to $(1 - 2\,\mathbb{E}[\tau])^{-2}$, indicating a room for improvement of the upper bound in the logistic and probit cases (the upper bound for the laplacian case we established is optimal).

Now let us investigate the conditions in Theorem 2. Note that we did not optimize the constants. For example, the condition $r \geq 16$ can be relaxed to, e.g., $r \geq 4$. The condition $\alpha^2 r d_2 \geq \mathrm{C}_0$ is easy to satisfy, especially in the high-dimensional regime where $r, d_2 \to \infty$. We also point out that it is very mild to assume $g'(x)$ is decreasing in $\mathbb{R}^+$. It amounts to imposing that the probability density function has a non-increasing tail, which holds for the popular statistical models in Section 3.1. Finally, when the rank $r \leq \mathcal{O}\left(d_1/\alpha^2\right)$, the right-hand side of the inequality in the theorem is always larger than $r\alpha^2/d_1$. It turns out that under the setting $\alpha = \Theta(1)$, the lower bound holds even when the matrix rank is of the same order of the dimension. We summarize the established bounds in Table 1.

3.3 Proof Sketch

We will consider the centralized loss function

$$\bar{L}_{\Omega,Y'}(X) := L_{\Omega,Y'}(X) - L_{\Omega,Y'}(0).$$

The following lemma is crucial for our analysis.

Lemma 3. *Let the set* $\mathcal{S}$ *be*

$$\mathcal{S} = \left\{X \in \mathbb{R}^{d_1 \times d_2} :\ \|X\|_* \leq \alpha\sqrt{r d_1 d_2}\right\}.$$

Write

$$G_{\Omega,Y'} = \sup_{X \in \mathcal{S}} \left|\bar{L}_{\Omega,Y'}(X) - \mathbb{E}\,\bar{L}_{\Omega,Y'}(X)\right|,$$
$$\bar{G} = \alpha\rho_\alpha^+ \sqrt{r}\sqrt{n(d_1 + d_2) + d_1 d_2 \log(d_1 d_2)}.$$

Then it follows that

$$\Pr\left(G_{\Omega,Y'} \geq \mathrm{C}_0 \bar{G}\right) \leq \frac{\mathrm{C}_1}{d_1 + d_2},$$

for some absolute constants C_0 *and* C_1.

Recall that the likelihood function we defined in (6) is not averaged by n. Hence, the above lemma suggests that when n is large enough, the shifted loss $\frac{1}{n}\bar{L}_{\Omega,Y'}(X)$ concentrates around its expectation with the rate $\mathcal{O}\left(1/\sqrt{n}\right)$.

On the other hand, by algebra we can show that

$$D\left(g(M)||g(\widehat{M})\right) \leq \frac{2}{n} G_{\Omega,Y'},$$

where the left-hand side is the KL divergence which can be further lower bounded

$$\frac{1}{d_1 d_2}\left\|\widehat{M} - M\right\|_F^2 \leq 8\rho_\alpha^+ D\left(g(M)||g(\widehat{M})\right).$$

This immediately implies Theorem 1 after some rearrangements.

The lower bound follows from standard information theoretic arguments. We construct a set of matrices that satisfy $(A2)$ and $(A3)$ but the discrepancy between the members of this set is large in terms of Frobenius norm. We then show that for any true matrix M coming from this set, it is not easy for any recovery algorithm to output a solution that is quite close to it. This suggests a lower bound as stated in Theorem 2. See Appendix B for the full proof.

4 Experiments

We complement our theoretical findings by performing a comprehensive set of experiments on both simulated data and realistic problems. In particular, for synthetic data our focus is on how the estimation error changes with the sample size n and the random sign flipping noise. For the real-world data, we will demonstrate that the proposed estimator works well even when it is fed with inaccurate information of the noise parameter τ. The solver for the convex program (5) is publicly available at Davenport's homepage, and we follow their default settings.

4.1 Simulation

We first elaborate the experimental settings.

Data. For simplicity, we set $d_1 = d_2 = d$ where $d = 200$. We randomly generate the true real-valued matrix $M \in \mathbb{R}^{d \times d}$ such that it has rank r and $\|M\|_\infty \leq 1$. To be more concrete, we construct two matrices $U, V \in \mathbb{R}^{d \times r}$ where the entries are drawn i.i.d. from a uniform distribution on the interval $[-1, 1]$. The low-rank matrix M is then given by the product $UV^\top$ followed by a normalization (that is, $M \leftarrow M/\|M\|_\infty$). Given a sample size n, the index set Ω is picked uniformly random such that $|\Omega| = n$. The noisy observation Y'_Ω depends on the choice of $f(x)$ and the flipping

Table 1: **Upper and lower bounds on the sample complexity in the regime where α is a constant.**

$f(x)$	Upper bound	Lower bound
Logistic	$(1-2\,\mathbb{E}[\tau])^{-4}\,r(d_1+d_2)\log(d_1 d_2)$	$(1-2\,\mathbb{E}[\tau])^{-2}\,r\max\{d_1,d_2\}$
Probit	$(1-2\,\mathbb{E}[\tau])^{-6}\,r(d_1+d_2)\log(d_1 d_2)$	$(1-2\,\mathbb{E}[\tau])^{-2}\,r\max\{d_1,d_2\}$
Laplacian	$(1-2\,\mathbb{E}[\tau])^{-2}\,r(d_1+d_2)\log(d_1 d_2)$	$(1-2\,\mathbb{E}[\tau])^{-2}\,r\max\{d_1,d_2\}$

parameter τ (see (3)). Here, we choose the probit regression for $f(x)$, i.e., $f(x)=\Phi(x/\sigma)$, the cumulative density function of zero-mean Gaussian distribution. The parameter $\sigma = 0.3$.

Evaluation. We measure the discrepancy between the recovered matrix $\widehat{M}$ and the true matrix M by the mean squared error (MSE). Each experiment to be showed are conducted for 5 trials, and we report the averaged MSE.

Our first empirical study focuses on the error curve against the sample size when τ is fixed as a scalar (so there is no randomness in τ). We point out that though τ is a deterministic quantity, the flipping noise is still random. Such a noise model was widely studied in the context of 1-bit compressed sensing [37, 45]. We set $\tau = 0.2$ which means for all $(i,j)\in\Omega$, the component Y_{ij} is flipped with probability 0.2. We plot the curves of MSE in Figure 1 where we also vary the rank r from 1 to 10. Note that a larger rank indicates a more complicated problem, hence we need to draw more observations to achieve a low error, as illustrated in this figure. Also note that in the right panel, the x-axis is $d/\sqrt{n}$ (n is the sample size), and we find that the statistical error scales approximately linear with it, which matches our theoretical prediction.

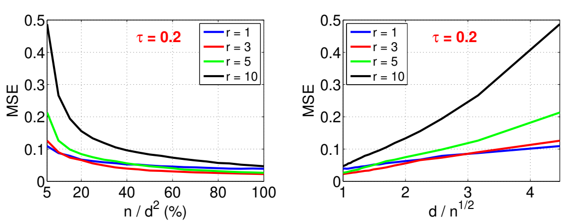

Figure 1: **Estimation error against sample size under fixed τ.** The x-axis is properly normalized by a constant for a better view. The statistical error is approximately linear with $1/\sqrt{n}$.

Then we fix the rank $r=3$, and tune the parameter τ from 0 to 0.4. Note that τ is still a deterministic quantity. For each value of τ, we plot the error curves in Figure 2. We observe that the recovery becomes challenging when the data are grossly corrupted. In the right panel, one can also observe a two-phase behavior in which the MSE varies approximately linearly with $d/\sqrt{n}$ when it is small, whereas when $d/\sqrt{n}$ is large the MSE for different values of τ more or less converge to a common curve. This is because in Figure 2 d is fixed, and a large $d/\sqrt{n}$ means a small sample size n. In this scenario, no algorithm is able to recover the true matrix (which is reflected by the large error).

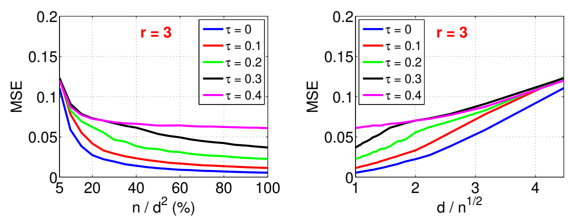

Figure 2: **Estimation error against sample size under fixed rank.**

Now we investigate the situation where τ itself is a random variable. A remarkable implication of our theoretical analysis is that the random variable τ affects the recovery only through its mean. We verify this by randomly generating 3 different distributions for τ, say $\mathcal{D}_1$, $\mathcal{D}_2$ and $\mathcal{D}_3$. For each distribution $\mathcal{D}_i$, τ takes value from $\{\tau_{ik}\}_{k=1}^4$ with corresponding probability $\{p_{ik}\}_{k=1}^4$. The configuration $\{\tau_{ik}, p_{ik}\}_{k=1}^4$ is generated randomly, but subject to the constraints that **(i)** each τ_{ik} lies in the interval $[0, 1/2)$; **(ii)** $\mathbb{E}[\tau]=0.2$; and **(iii)** $\sum_{k=1}^4 p_{ik} = 1$ for a given i. Then for each distribution $\mathcal{D}_i$, we manually corrupt the clean data Y_Ω and run the solver to obtain an estimate. The results are recorded in Figure 3 where we use the logarithmic scale for the y-axis to magnify the difference for the curves of these distributions. Even by doing so, we find that the three curves are almost lying on top of each other, which verifies our theoretical finding that the statistical error only depends on the mean of τ.

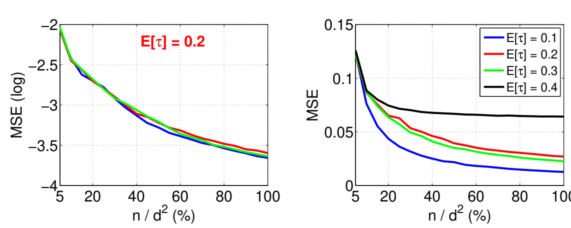

Figure 3: **Estimation error against sample size under the same and different noise expectation.**

Finally, we generate 4 distributions with different $\mathbb{E}[\tau]$ using the same scheme as above and illustrate the results in Figure 3. These curves again show that the flipping noise poses challenges for accurate recovery.

4.2 MovieLens-100K under Bounded Noise

Data. As we have verified the theoretical findings via simulation, now we are moving on to conduct experiments on real-world data set: MovieLens-100K. It consists of 100,000 ratings collected from 1000 users on 1700 movies, where the original rating is an integer from 1 to 5. We threshold the ratings by their average to obtain a binary matrix. Then we randomly choose 95% for training and the remaining for testing.

Noise Model and Our Estimator. Our goal is to examine if the proposed estimator is resilient to inaccurate information of the noise parameter τ. This is crucial for practical applications since the noise is usually a hidden variable and the best we can hope is to know its upper bound. Thus, we corrupt each clean entry y_{ij} by flipping it with arbitrary probability τ_{ij}. Note that this is a much more challenging case than before in that each τ_{ij} may follow a different distribution. We assume that all $\tau_{ij} \leq \tau_{\max} < 1/2$ and we have the knowledge of $\tau_{\max}$, which is exactly the bounded noise model [32]. Our estimator (5), specifically the function $g(x)$ is constructed with $\tau_{\max}$ in place of $\mathbb{E}[\tau]$.

Evaluation. We consider the recovery percentage of the sign patterns, with the purpose of simulating the 0/1 recommender systems. We use the logistic function for $f(x)$. Thus, when the program (5) returns the estimate $\widehat{M}$, we threshold $f(\widehat{M})$ by 0.5 entry-wisely to obtain a binary estimation $\widehat{Y}$.

We first illustrate that it is safe to apply our method even when there is no noise in the data. Note that [17] was developed specifically for the case. In this scenario, we train two models: one with the knowledge that $\tau = 0$ (which is exactly the model of [17]) and another one with $\tau_{\max} = 0.4$. The results are recorded in the first row of Table 2. Interestingly, we observe that our estimator with $\tau_{\max}$ performs slightly better than the ones with the knowledge of noise. Their downgrade is probably owing to the intrinsic noise in the data, e.g., noisy inputs from the users.

Table 2: **Accuracy (%) of sign recovery.**

	[17]	Ours with $\mathbb{E}[\tau]$	Ours with $\tau_{\max}$
Noise-free	74.9	74.9	75.0
Bounded	70.3	72.5	72.3

Then we consider the bounded noise model where each observed entry is flipped with an arbitrary probability smaller than 0.2. We train the model of [17] with the noisy data, and train two models based on our robust formulation (5): one with the exact estimate of $\mathbb{E}[\tau]$ which is 0.1 and another with a rough estimate of $\tau_{\max} = 0.4$. From the second row of Table 2, it is not hard to see that [17] degrades a lot since it is noise-oblivious while our models performs well. In particular, we find that even we do not know the exact information of the noise, it is still possible to achieve comparable performance.

Finally, we tune the expected value of τ from 0 to 0.3 with a step size 0.05 (so totally 7 values of $\mathbb{E}[\tau]$), but we alway set $\tau_{\max} = 0.4$. We train our two models as before, and compare the curves produced by the one with the knowledge of $\mathbb{E}[\tau]$ and the one with only $\tau_{\max}$. The results are plotted in Figure 4. As we expected, our estimator works well under the bounded noise model, and is always superior to [17]. This again demonstrates the robustness of our model.

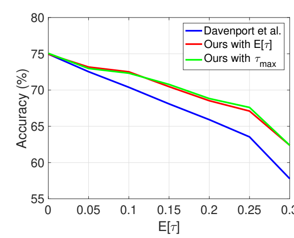

Figure 4: **Recovery accuracy of sign patterns under bounded noise.**

5 Conclusion

In this paper, we have introduced the noisy 1-bit matrix completion model, where each observed entry is flipped with some probability controlled by a random variable $\tau \in [0, 1/2)$. It has been shown that under rather mild conditions on the sampling scheme and the true matrix, a simple maximum likelihood estimator guarantees accurate recovery with high probability. Along with our analysis, we have established that the random variable τ enters the sample complexity only through its mean. When the binary data are generated from a Laplacian distribution, we have demonstrated that the upper bound matches the lower bound (up to a logarithmic factor). We have carried out a variety of numerical study to show that our theorems match perfectly the empirical results. Perhaps somewhat surprisingly, we also demonstrate on a benchmark data set that the proposed estimator is resilient to bounded noise. We believe it is a promising direction to pursue a more refined analysis under such noise model, and to conduct experiments on larger data sets.

Acknowledgements

J. Shen is supported by the startup funding from Stevens Institute of Technology. We thank the anonymous reviewers for pointing out several related work and for suggestions on improving the clarity.

References

[1] Pranjal Awasthi, Maria-Florina Balcan, Nika Haghtalab, and Ruth Urner. Efficient learning of linear separators under bounded noise. In *Proceedings of the 28th Conference on Learning Theory*, pages 167–190, 2015.

[2] Pranjal Awasthi, Maria-Florina Balcan, Nika Haghtalab, and Hongyang Zhang. Learning and 1-bit compressed sensing under asymmetric noise. In *Proceedings of the 29th Conference on Learning Theory*, pages 152–192, 2016.

[3] Pranjal Awasthi, Maria-Florina Balcan, and Philip M. Long. The power of localization for efficiently learning linear separators with noise. *Journal of the ACM*, 63(6):50:1–50:27, 2017.

[4] Maria-Florina Balcan and Hongyang Zhang. Noise-tolerant life-long matrix completion via adaptive sampling. In *Proceedings of the 30th Annual Conference on Neural Information Processing Systems*, pages 2955–2963, 2016.

[5] Sonia A. Bhaskar. Probabilistic low-rank matrix completion from quantized measurements. *Journal of Machine Learning Research*, 17(60):1–34, 2016.

[6] Thomas Blumensath and Mike E. Davies. Iterative hard thresholding for compressed sensing. *Applied and Computational Harmonic Analysis*, 27(3):265–274, 2009.

[7] Petros Boufounos and Richard G. Baraniuk. 1-bit compressive sensing. In *Proceedings of the 42nd Annual Conference on Information Sciences and Systems*, pages 16–21, 2008.

[8] Emmanuel J. Candès, Xiaodong Li, Yi Ma, and John Wright. Robust principal component analysis? *Journal of the ACM*, 58(3):11:1–11:37, 2011.

[9] Emmanuel J. Candès and Yaniv Plan. Matrix completion with noise. *Proceedings of the IEEE*, 98(6):925–936, 2010.

[10] Emmanuel J. Candès and Benjamin Recht. Exact matrix completion via convex optimization. *Foundations of Computational Mathematics*, 9(6):717–772, 2009.

[11] Emmanuel J. Candès and Terence Tao. Decoding by linear programming. *IEEE Transactions on Information Theory*, 51(12):4203–4215, 2005.

[12] Emmanuel J. Candès and Terence Tao. The power of convex relaxation: near-optimal matrix completion. *IEEE Transactions on Information Theory*, 56(5):2053–2080, 2010.

[13] Venkat Chandrasekaran, Benjamin Recht, Pablo A. Parrilo, and Alan S. Willsky. The convex geometry of linear inverse problems. *Foundations of Computational Mathematics*, 12(6):805–849, 2012.

[14] Scott Shaobing Chen, David L. Donoho, and Michael A. Saunders. Atomic decomposition by basis pursuit. *SIAM Journal on Scientific Computing*, 20(1):33–61, 1998.

[15] Yudong Chen. Incoherence-optimal matrix completion. *IEEE Transactions on Information Theory*, 61(5):2909–2923, 2015.

[16] Yudong Chen, Huan Xu, Constantine Caramanis, and Sujay Sanghavi. Matrix completion with column manipulation: Near-optimal sample-robustness-rank tradeoffs. *IEEE Transactions on Information Theory*, 62(1):503–526, 2016.

[17] Mark A. Davenport, Yaniv Plan, Ewout van den Berg, and Mary Wootters. 1-bit matrix completion. *Information and Inference: A Journal of the IMA*, 3(3):189–223, 2014.

[18] David L. Donoho. Compressed sensing. *IEEE Transactions on Information Theory*, 52(4):1289–1306, 2006.

[19] Maryam Fazel, Haitham Hindi, and Stephen P. Boyd. A rank minimization heuristic with application to minimum order system approximation. In *Proceedings of the American Control Conference*, volume 6, pages 4734–4739, 2001.

[20] Sivakant Gopi, Praneeth Netrapalli, Prateek Jain, and Aditya V. Nori. One-bit compressed sensing: Provable support and vector recovery. In *Proceedings of the 30th International Conference on Machine Learning*, pages 154–162, 2013.

[21] David Gross. Recovering low-rank matrices from few coefficients in any basis. *IEEE Transactions on Information Theory*, 57(3):1548–1566, 2011.

[22] Ankit Gupta, Robert D. Nowak, and Benjamin Recht. Sample complexity for 1-bit compressed sensing and sparse classification. In *Proceedings of the IEEE International Symposium on Information Theory*, pages 1553–1557, 2010.

[23] Laurent Jacques, Jason N. Laska, Petros T. Boufounos, and Richard G. Baraniuk. Robust 1-bit compressive sensing via binary stable embeddings of sparse vectors. *IEEE Transactions on Information Theory*, 59(4):2082–2102, 2013.

[24] Prateek Jain, Praneeth Netrapalli, and Sujay Sanghavi. Low-rank matrix completion using alternating minimization. In *Proceedings of the 45th Annual Symposium on the Theory of Computing*, pages 665–674, 2013.

[25] Chi Jin, Sham M. Kakade, and Praneeth Netrapalli. Provable efficient online matrix completion via non-convex stochastic gradient descent. In *Proceedings of the 30th Annual Conference on Neural Information Processing Systems*, pages 4520–4528, 2016.

[26] Raghunandan H. Keshavan, Andrea Montanari, and Sewoong Oh. Matrix completion from a few entries. *IEEE Transactions on Information Theory*, 56(6):2980–2998, 2010.

[27] Raghunandan H. Keshavan, Andrea Montanari, and Sewoong Oh. Matrix completion from noisy entries. *Journal of Machine Learning Research*, 11:2057–2078, 2010.

[28] Karin Knudson, Rayan Saab, and Rachel Ward. One-bit compressive sensing with norm estimation. *IEEE Transactions on Information Theory*, 62(5):2748–2758, 2016.

[29] Jean Lafond. Low rank matrix completion with exponential family noise. In *Proceedings of the 28th Conference on Learning Theory*, pages 1224–1243, 2015.

[30] Michel Ledoux and Michel Talagrand. *Probability in Banach Spaces: Isoperimetry and Processes*. Springer-Verlag Berlin Heidelberg, 1991.

[31] Yu Lu and Sahand N. Negahban. Individualized rank aggregation using nuclear norm regularization. In *Proceedings of the 53rd Annual Allerton Conference on Communication, Control, and Computing*, pages 1473–1479, 2015.

[32] Pascal Massart and Élodie Nédélec. Risk bounds for statistical learning. *The Annals of Statistics*, pages 2326–2366, 2006.

[33] Sahand Negahban and Martin J. Wainwright. Restricted strong convexity and weighted matrix completion: Optimal bounds with noise. *Journal of Machine Learning Research*, 13:1665–1697, 2012.

[34] Dohyung Park, Joe Neeman, Jin Zhang, Sujay Sanghavi, and Inderjit S. Dhillon. Preference completion: Large-scale collaborative ranking from pairwise comparisons. In *Proceedings of the 32nd International Conference on Machine Learning*, pages 1907–1916, 2015.

[35] Yagyensh C. Pati, Ramin Rezaiifar, and Perinkulam S. Krishnaprasad. Orthogonal matching pursuit: Recursive function approximation with applications to wavelet decomposition. In *Conference Record of The 27th Asilomar Conference on Signals, Systems and Computers*, pages 40–44, 1993.

[36] Yaniv Plan and Roman Vershynin. One-bit compressed sensing by linear programming. *Communications on Pure and Applied Mathematics*, 66(8):1275–1297, 2013.

[37] Yaniv Plan and Roman Vershynin. Robust 1-bit compressed sensing and sparse logistic regression: A convex programming approach. *IEEE Transactions on Information Theory*, 59(1):482–494, 2013.

[38] Yaniv Plan, Roman Vershynin, and Elena Yudovina. High-dimensional estimation with geometric constraints. *Information and Inference: A Journal of the IMA*, 6(1):1–40, 2017.

[39] Benjamin Recht, Maryam Fazel, and Pablo A. Parrilo. Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. *SIAM Review*, 52(3):471–501, 2010.

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

[41] Jie Shen and Ping Li. On the iteration complexity of support recovery via hard thresholding pursuit. In *Proceedings of the 34th International Conference on Machine Learning*, pages 3115–3124, 2017.

[42] Jie Shen and Ping Li. Partial hard thresholding: Towards a principled analysis of support recovery. In *Proceedings of the 31st Annual Conference on Neural Information Processing Systems*, pages 3127–3137, 2017.

[43] Jie Shen and Ping Li. A tight bound of hard thresholding. *Journal of Machine Learning Research*, 18(208):1–42, 2018.

[44] Jie Shen, Huan Xu, and Ping Li. Online optimization for max-norm regularization. In *Proceedings of the 28th Annual Conference on Neural Information Processing Systems*, pages 1718–1726, 2014.

[45] Martin Slawski and Ping Li. Linear signal recovery from b-bit-quantized linear measurements: precise analysis of the trade-off between bit depth and number of measurements. *CoRR*, abs/1607.02649, 2016.

[46] Nathan Srebro, Jason D. M. Rennie, and Tommi S. Jaakkola. Maximum-margin matrix factorization. In *Proceedings of the 18th Annual Conference on Neural Information Processing Systems*, pages 1329–1336, 2004.

[47] Robert Tibshirani. Regression shrinkage and selection via the Lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1):267–288, 1996.

[48] Joel A. Tropp and Anna C. Gilbert. Signal recovery from random measurements via orthogonal matching pursuit. *IEEE Transactions on Information Theory*, 53(12):4655–4666, 2007.

[49] Joel A. Tropp and Stephen J. Wright. Computational methods for sparse solution of linear inverse problems. *Proceedings of the IEEE*, 98(6):948–958, 2010.

[50] Martin J. Wainwright. Sharp thresholds for high-dimensional and noisy sparsity recovery using ℓ_1-constrained quadratic programming (Lasso). *IEEE Transactions on Information Theory*, 55(5):2183–2202, 2009.

[51] Jing Wang, Jie Shen, and Ping Li. Provable variable selection for streaming features. In *Proceedings of the 35th International Conference on Machine Learning*, pages 5158–5166, 2018.

[52] Jing Wang, Jie Shen, Ping Li, and Huan Xu. Online matrix completion for signed link prediction. In *Proceedings of the 10th ACM International Conference on Web Search and Data Mining*, pages 475–484, 2017.

[53] Songbai Yan and Chicheng Zhang. Revisiting perceptron: Efficient and label-optimal learning of halfspaces. In *Proceedings of the 31st Annual Conference on Neural Information Processing Systems*, pages 1056–1066, 2017.