MATRIX RECOVERY FROM QUANTIZED AND CORRUPTED MEASUREMENTS

Andrew S. Lan[1], Christoph Studer[2], and Richard G. Baraniuk[1]

[1]Rice University; e-mail: {sl29, richb}@rice.edu
[2]Cornell University; e-mail: studer@cornell.edu

ABSTRACT

This paper deals with the recovery of an unknown, low-rank matrix from quantized and (possibly) corrupted measurements of a subset of its entries. We develop statistical models and corresponding (multi-)convex optimization algorithms for quantized matrix completion (Q-MC) and quantized robust principal component analysis (Q-RPCA). In order to take into account the quantized nature of the available data, we jointly learn the underlying quantization bin boundaries and recover the low-rank matrix, while removing potential (sparse) corruptions. Experimental results on synthetic and two real-world collaborative filtering datasets demonstrate that directly operating with the quantized measurements—rather than treating them as real values—results in (often significantly) lower recovery error if the number of quantization bins is less than about 10.

Index Terms— Quantization, convex optimization, matrix completion, robust principal component analysis.

1. INTRODUCTION

1.1. Matrix completion and robust PCA

Matrix completion (MC) aims at recovering an unknown, low-rank matrix $\mathbf{L} \in \mathbb{R}^{m \times n}$ from a subset of real-valued measurements $Y_{i,j} \in \mathbb{R}$ with $(i, j) \in \Omega_{\text{obs}}$, where Ω_{obs} is the set of indices of the observed data. Corresponding theoretical results for MC have shown that only $O(r\,n\,\text{polylog}(n))$ randomly observed entries in Ω_{obs} are required to exactly recover the $n \times n$, rank-r matrix $\mathbf{L}$ (see, e.g., [1]). An extension of MC that is capable of recovering low-rank matrices from a set of sparsely corrupted measurements, known as robust principal component analysis (RPCA) [2, 3], further assumes that the unknown matrix $\mathbf{X} \in \mathbb{R}^{m \times n}$ to be recovered is constructed as the sum of a low-rank matrix $\mathbf{L}$ and a sparse matrix $\mathbf{S}$; this extension renders MC resilient to outliers in the observed entries. The recent popularity of MC and RPCA is mainly due to the fact that they find use in a large number of practical applications and fields, including (but not limited to) collaborative filtering [4, 5], foreground-background separation [2], financial data analysis [6], and compressive sensing [7].

1.2. 1-bit matrix completion

In many practical applications, the common assumption of MC that the measurements $Y_{i,j}$ are real-valued is *not valid*. In collaborative filtering, for example, the observed data corresponds to user ratings of items (e.g., movies), which are from a discrete set of values, such as binary-valued (1-bit) [8, 9] or quantized[1] [10–12], e.g., from the discrete and ordered set $\{1, \ldots, 5\}$. Although one can apply existing MC and RPCA algorithms to quantized measurements by treating them as real values (e.g., by mapping them onto the set of real numbers), recent results in 1-bit MC [13] have shown that directly using the available binary-valued data in combination with a suitable measurement model often results in substantially better recovery performance [11, 12]. This line of research has shown that only $O(r\,n)$ randomly-observed and binary-valued entries are sufficient to obtain a robust estimate of the $n \times n$, rank-r matrix $\mathbf{L}$ [13].

Unfortunately, extending the algorithms and recovery guarantees for 1-bit MC [11–13] to the general case of quantized and (possibly) corrupted measurements is not straightforward. Hence, new models and corresponding efficient MC and RPCA algorithms are of paramount importance for applications relying on quantized and (possibly) corrupted measurements, such as in recommender systems [4, 14], educational assessments [12], and psychological tests [15].

1.3. Contributions

In this paper, we develop new statistical models for quantized and (possibly) corrupted measurements of a low-rank matrix $\mathbf{L}$, and two corresponding (multi-)convex optimization-based recovery algorithms, namely *quantized matrix completion* (Q-MC) and *quantized robust principal component analysis* (Q-RPCA). We base our algorithms on the FISTA framework [16] in order to solve the optimization problems with low computational complexity. In addition, our algorithms directly estimate the set of quantization bin boundaries from observed data using an alternating optimization procedure, since knowledge of the quantization bin boundaries is, in general, unavailable in practice. We demonstrate the effectiveness of our methods using experiments on synthetic and two real-world collaborative filtering datasets.

---

[1]The machine learning literature commonly uses the term "ordinal data" for quantized measurements; we will use both terms interchangeably.

2. STATISTICAL MODELS FOR QUANTIZED MATRIX RECOVERY

Q-MC and Q-RPCA aim at recovering an unknown, low-rank matrix $\mathbf{L}$ from a subset of quantized and (possibly) corrupted measurements. In the Q-MC case, the matrix $\mathbf{X}$ to be recovered is assumed to be low-rank, i.e., $\mathbf{X} = \mathbf{L}$; in the Q-RPCA case, the matrix $\mathbf{X}$ to be recovered is composed of a low-rank and a sparse corruption matrix, i.e., $\mathbf{X} = \mathbf{L}+\mathbf{S}$. In both cases, we assume that the matrix $\mathbf{X}$ of interest is of dimension $m \times n$. In collaborative filtering applications, the dimension m often refers to the number of *items* and n to the number of *users*.

2.1. Quantization model

Let $Y_{i,j}$ represent the quantized measurement of the $(i,j)^{\text{th}}$ entry of the (unknown) matrix $\mathbf{X}$ (or the rating of the j^{th} user on the i^{th} item), which are from a set of P ordered labels,[2] i.e., $Y_{i,j} \in \mathcal{O}$ with $\mathcal{O} = \{1, \ldots, P\}$. We propose the following model for the quantized measurements $Y_{i,j}$:

$$Y_{i,j} = \mathcal{Q}(X_{i,j} + \epsilon_{i,j}), \ (i,j) \in \Omega_{\text{obs}},$$
$$\epsilon_{i,j} \sim \text{Logistic}(0,1) \quad \text{or} \quad \epsilon_{i,j} \sim \mathcal{N}(0,1). \quad (1)$$

The quantity $\epsilon_{i,j}$ models the uncertainty on each measurement of $X_{i,j}$. $\text{Logistic}(0,1)$ denotes a logistic distribution with zero mean and unit scale; $\mathcal{N}(0,1)$ denotes the standard normal distribution. The set $\Omega_{\text{obs}} \subseteq \{1, \ldots, m\} \times \{1, \ldots, n\}$ contains the indices associated to the observed measurements $Y_{i,j}$. In (1), the function $\mathcal{Q}(\cdot) : \mathbb{R} \to \mathcal{O}$ corresponds to a scalar quantizer that maps a real number to one of the P ordered labels according to

$$\mathcal{Q}(x) = p \quad \text{if } \omega_{p-1} < x \leq \omega_p, \ p \in \mathcal{O},$$

where $\{\omega_0, \ldots, \omega_P\}$ represent the set of quantization bin boundaries satisfying $\omega_0 \leq \omega_1 \leq \cdots \leq \omega_{P-1} \leq \omega_P$, with ω_0 and ω_P denoting the lower and upper quantization bin boundaries of the quantizer $\mathcal{Q}(\cdot)$. Note that in most situations, we have $\omega_0 = -\infty$ and $\omega_P = \infty$.

We emphasize that the quantization bin boundaries are, in general, unknown, and depend on the application at hand. For example, in recommender systems, the ratings 4 and 5 (out of 5) often stand for "good" and "excellent," which can be interpreted very differently by different users. We therefore propose an automatic technique to learn the quantization bin boundaries directly from the available data (see Sec. 3). For notational simplicity, we assume that all measurements $Y_{i,j}$ share a common set of quantization bin boundaries. In practice, however, a set of unique bin boundaries can be defined either for each item (denoted as $\omega_0^i \leq \cdots \leq \omega_P^i, i = 1, \ldots, m$), associated with each row of $\mathbf{X}$, or for each user (denoted as $\omega_0^j \leq \cdots \leq \omega_P^j, j = 1, \ldots, n$), associated with each column of $\mathbf{X}$. A comparison of different assumptions on the quantization bin boundaries will be provided in Sec. 4.2.

2.2. Statistical model

The measurement model in (1) leads to the following statistical input–output relation for each observation $Y_{i,j}$:

$$p(Y_{i,j} = p \mid X_{i,j}) = \Phi(\omega_p - X_{i,j}) - \Phi(\omega_{p-1} - X_{i,j}), \quad (2)$$

where $\Phi(x)$ corresponds to an inverse link function. For logistic noise, we use $\Phi_{\text{log}}(x) = \frac{1}{1+e^{-x}}$, the inverse logit link function; for standard normal noise, we use $\Phi_{\text{pro}}(x) = \int_{-\infty}^{x} \mathcal{N}(s|0,1)\mathrm{d}s$, the inverse probit link function [17]. The Q-MC and Q-RPCA algorithms proposed in Sec. 3 can be formulated for both noise models. Note that we can rewrite (1) and (2) in compact form as

$$p(Y_{i,j} \mid X_{i,j}) = \Phi(U_{i,j} - X_{i,j}) - \Phi(W_{i,j} - X_{i,j}),$$

where the $m \times n$ matrices $\mathbf{U}$ and $\mathbf{W}$ contain the upper and lower bin boundaries corresponding to the measurements $Y_{i,j}$, i.e., we have $U_{i,j} = \omega_{Y_{i,j}}$ and $W_{i,j} = \omega_{Y_{i,j}-1}$.

We emphasize that the model (1) generalizes the common 1-bit MC model proposed in [13] to quantized data. In particular, 1-bit MC is a special case of our model with $P = 2$, $\{\omega_0, \omega_1, \omega_2\} = \{-\infty, 0, \infty\}$, and $\mathbf{X} = \mathbf{L}$. We next propose two algorithms that jointly learn the quantization bin boundaries and recover $\mathbf{L}$ (and also $\mathbf{S}$ in the case of Q-RPCA[3]), given the subset of quantized measurements $Y_{i,j}, (i,j) \in \Omega_{\text{obs}}$, under the assumptions that $\text{rank}(\mathbf{L}) \ll m, n$ and $\mathbf{S}$ is sparse (i.e., contains only a small number of non-zero entries).

3. ALGORITHMS FOR QUANTIZED MATRIX RECOVERY

3.1. Q-MC: Quantized matrix completion

In order to recover the low-rank matrix $\mathbf{L}$ from quantized measurements, we minimize the negative log-likelihood of $Y_{i,j}, (i,j) \in \Omega_{\text{obs}}$, given by (2), subject to a low-rank promoting constraint on $\mathbf{L}$. In particular, we seek to solve the following constrained optimization problem:

$$(\text{Q-MC}) \begin{cases} \underset{\mathbf{L}, \omega_0 \leq \cdots \leq \omega_P}{\text{minimize}} & -\sum_{i,j:(i,j)\in\Omega_{\text{obs}}} \log p(Y_{i,j}|L_{i,j}) \\ \text{subject to} & \|\mathbf{L}\|_* \leq \lambda. \end{cases}$$

Here, the nuclear norm constraint $\|\mathbf{L}\|_* \leq \lambda$, which is a convex relaxation of the low-rank constraint $\text{rank}(\mathbf{L}) \leq r$, promotes low-rankness of $\mathbf{L}$ [18]; the parameter $\lambda > 0$ is used to tune the rank of $\mathbf{L}$. In practice, we select the parameter λ via cross-validation or AIC/BIC [17]. We note that one can also use a Lagrange multiplier [19] to raise the constraint $\|\mathbf{L}\|_* \leq \lambda$ of (Q-MC) to the objective function as a regularizer, which is equivalent to the common approach for (real-valued) MC [20]. Experimental results suggest that selecting the optimal value of the parameter λ in the constraint is—in

---

[2] For the sake of simplicity, we only discuss the case where all observations have the same number of quantization bins P; the generalization to the case of a different number of bins for each measurement is straightforward.

[3] Note that for Q-RPCA, it is not possible to recover the sparse matrix $\mathbf{S}$ if only a subset of the entries of $\mathbf{X} = \mathbf{L} + \mathbf{S}$ are observed—just like RPCA.

most cases—simpler (i.e., less sensitive to small variations) than selecting the optimal value of λ in the regularizer.

Since the negative log-likelihood of the logit and probit functions are convex in $\mathbf{L}$ when keeping the quantization bin boundaries $\omega_0, \ldots, \omega_P$ fixed, and (Q-MC) is convex in each individual quantization bin boundary while holding the other ones and $\mathbf{L}$ fixed [17], the problem (Q-MC) is *multi-convex* in all the variables $\mathbf{L}$ and $\omega_0, \ldots, \omega_P$. Therefore, (Q-MC) can be solved approximately via block-coordinate descent [21].

The Q-MC algorithm consists of an inner and an outer iteration. We first initialize the matrix $\mathbf{L}$ with an all-zero matrix and the quantization bin boundaries with $\omega_0 = \Phi^{-1}(0)$, $\omega_1 = \Phi^{-1}(\frac{1}{P}), \ldots, \omega_P = \Phi^{-1}(1)$. Then, in each outer iteration, the Q-MC algorithm performs two steps consecutively: (i) We hold all quantization bin boundaries fixed and optimize $\mathbf{L}$, which can be solved using an iterative procedure detailed below (forming the inner iterations); (ii) We hold $\mathbf{L}$ fixed and optimize the bin boundaries one by one, while holding the others fixed. This step can, for example, be accomplished using the secant method [22]. The associated details are omitted due to space constraints. We terminate the algorithm either if a maximum number of outer iterations is reached or if the decrease in the objective function between consecutive outer iterations is smaller than a predefined threshold.

The problem of optimizing $\mathbf{L}$ while holding the quantization bin boundaries fixed corresponds to the following nuclear-norm constrained logit/probit regression problem:

$$\begin{cases} \underset{\mathbf{L}}{\text{minimize}} & f(\mathbf{L}) = -\sum_{i,j:\,(i,j)\in\Omega_{\text{obs}}} \log p(Y_{i,j}|L_{i,j}) \\ \text{subject to} & \|\mathbf{L}\|_* \leq \lambda, \end{cases}$$

which can be solved efficiently via the FISTA framework [16] for non-smooth constrained/regularized optimization problems. Starting with an initialization of the matrix $\mathbf{L}$, at each inner iteration $\ell = 1, 2, \ldots$, the algorithm performs a gradient step to reduce the cost function followed by a projection step to make the solution satisfy the non-smooth nuclear-norm constraint. Both steps are repeated until a maximum number of inner iterations K_{max} is reached or the change in $\mathbf{L}$ between consecutive iterations is below a given threshold.

The *gradient step* aims at reducing the cost function $f(\mathbf{L})$, and is given by $\widehat{\mathbf{L}}^{\ell+1} \leftarrow \mathbf{L}^\ell - s_\ell \nabla f$. Here, s_ℓ is the step-size at iteration ℓ. For simplicity, we use a constant step-size $s_\ell = 1/L$, where L is the Lipschitz constant, which is given by $L_{\text{log}} = 1/4$ for the inverse logit link and $L_{\text{pro}} = 1$ for the inverse probit link. The gradient of the negative log-likelihood function of (Q-MC) with respect to $\mathbf{L}$ is given by

$$[\nabla f]_{i,j} = \begin{cases} \frac{\Phi'(W_{i,j}-L_{i,j})-\Phi'(U_{i,j}-L_{i,j})}{\Phi(U_{i,j}-L_{i,j})-\Phi(W_{i,j}-L_{i,j})} & \text{if } (i,j) \in \Omega_{\text{obs}}, \\ 0 & \text{otherwise.} \end{cases}$$

The derivatives of the inverse link function $\Phi'(x)$ can be calculated as $\Phi'_{\text{log}}(x) = \frac{1}{2+e^{-x}+e^x}$ and $\Phi'_{\text{pro}}(x) = \mathcal{N}(x\,|\,0,1)$.

The *projection step* aims to impose low-rankness on $\mathbf{L}$. This step is given by a projection onto the nuclear norm ball with radius λ, which is equivalent to taking the singular value

decomposition (SVD) of $\mathbf{L}$ followed by a projection of the singular-value vector onto the ℓ_1 ball with radius λ (see [20] for the details). Thus, the projection step is given by

$$\mathbf{L}^{\ell+1} \leftarrow \widetilde{\mathbf{U}}\text{diag}(\mathbf{s})\widetilde{\mathbf{V}}^T, \text{ with } \mathbf{s} = \text{P}_\lambda(\text{diag}(\mathbf{S})), \quad (3)$$

where $\widetilde{\mathbf{U}}\mathbf{S}\widetilde{\mathbf{V}}^T$ denotes the SVD of $\widehat{\mathbf{L}}^{\ell+1}$, and $\text{P}_\lambda(\cdot)$ the projection onto the ℓ_1-ball with radius λ, which can be computed efficiently using the algorithm detailed in [23].

3.2. Q-RPCA: Quantized robust PCA

The MC framework is able to handle uniform measurement noise but is prone to fail when parts of the data is corrupted by outliers. Real-valued robust principal component analysis (RPCA) [2] extends MC with the assumption that the matrix $\mathbf{X}$ to be recovered consist of a low-rank part $\mathbf{L}$ and a sparse corruption part $\mathbf{S}$; this modification renders RPCA resilient to outliers (and possibly leads to solutions $\mathbf{L}$ of lower rank than MC). In order to improve the robustness of Q-MC to outliers, we now extend the Q-MC algorithm detailed in Sec. 3.1 to quantized robust PCA (Q-RPCA). In particular, we seek to solve the following constrained optimization problem:

$$(\text{Q-RPCA}) \begin{cases} \underset{\mathbf{L},\mathbf{S},\omega_0 \leq \cdots \leq \omega_P}{\text{minimize}} & -\sum_{i,j:(i,j)\in\Omega_{\text{obs}}} \log p(Y_{i,j}|L_{i,j}+S_{i,j}) \\ \text{subject to} & \|\mathbf{L}\|_* \leq \lambda \text{ and } \|\mathbf{S}\|_1 \leq \sigma. \end{cases}$$

Here, $\|\mathbf{S}\|_1 \leq \sigma$ corresponds to an ℓ_1-norm constraint on the vectorized version of $\mathbf{S}$ that promotes sparsity among all entries of $\mathbf{S}$. In practice, we select the parameters λ and σ via cross-validation or AIC/BIC [17].

The (Q-RPCA) problem can be solved using a similar block coordinate descent procedure as in Sec. 3.1. In each outer iteration, we perform three steps: (i) hold the sparse part $\mathbf{S}$ and the quantization bin boundaries fixed and optimize for the low-rank part $\mathbf{L}$, (ii) hold $\mathbf{L}$ and the quantization bin boundaries fixed and optimize the sparse part $\mathbf{S}$, and (iii) hold both $\mathbf{L}$ and $\mathbf{S}$ fixed and optimize the quantization bin boundaries. The problem of optimizing the sparse matrix $\mathbf{S}$ while holding the other variables fixed is solved using a similar approach as in Sec. 3.1. Specifically, we use a projection step as in [23], i.e., $\mathbf{S}^{\ell+1} \leftarrow \text{P}_\sigma(\widehat{\mathbf{S}}^{\ell+1})$. Here, the ℓ_1-norm ball projection $\text{P}_\sigma(\cdot)$ operates on the vectorized version of $\widehat{\mathbf{S}}^{\ell+1}$.

4. EXPERIMENTS

We now validate the efficacy of Q-MC and Q-RPCA on both synthetic and real-world data. In all experiments, the parameters λ and σ are selected using cross-validation. All results are averaged over 25 independent Monte–Carlo trials.

4.1. Synthetic data

We start by evaluating the performance of Q-MC using synthetic data. Only the probit version of Q-MC will be used for
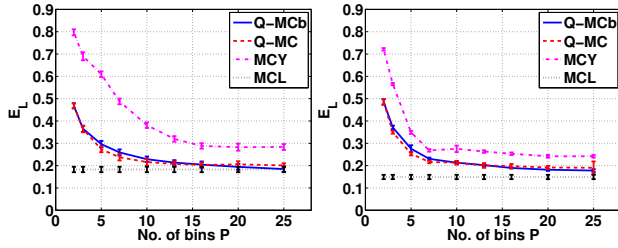
Fig. 1. Comparison of Q-MC versus real-valued MC for rank($\mathbf{L}$) = 5 and rank($\mathbf{L}$) = 10 in terms of the normalized recovery error $E_{\mathbf{L}}$.

		MovieLens		Dating Agency	
		MAE	RMSE	MAE	RMSE
	one set bin	0.7553	0.9443	1.1763	1.6655
Q-MC	user bin	0.7393	0.9301	1.1830	1.6700
	item bin	0.7394	0.9305	1.2026	1.6812
	one set bin	0.7562	0.9447	**1.1622**	**1.6570**
Q-RPCA	user bin	**0.7389**	**0.9294**	1.1818	1.6703
	item bin	0.7614	0.9467	1.1919	1.6749
MC		0.8588	1.1012	1.1784	1.6878
RPCA		0.8596	1.1042	1.1888	1.7230

Table 1. Performance comparison of Q-MC and Q-RPCA vs. real-valued MC and RPCA on predicting unobserved ratings for the MovieLens [24] and Dating Agency [14] datasets.

this experiment, since real-valued MC and RPCA typically assume a measurement model including Gaussian noise. We generate the synthetic data as follows. We set $m = n = 100$ and rank($\mathbf{L}$) to either 5 or 10. For each trial we generate $\mathbf{L}$ as the outer product of two low-dimensional random Gaussian matrices (according to the specified rank). The number of quantization bins is $P \in \{2, 3, 5, 7, 10, 13, 16, 20, 25\}$, and the bin boundaries $\omega_0, \ldots, \omega_P$ are chosen such that the number of entries in $\mathbf{L}$ that fall in each bin are roughly equal. The measurements $Y_{i,j}$ are then generated via (1), with 20% unobserved entries (chosen uniformly at random).

To assess the recovery performance of our proposed algorithms, we use the normalized recovery error defined as $E_{\mathbf{L}} = \|\mathbf{L} - \widehat{\mathbf{L}}\|_F^2 / \|\mathbf{L}\|_F^2$, where $\widehat{\mathbf{L}}$ is a re-scaled version of the recovered low-rank matrix having the same Frobenius norm as the ground-truth matrix $\mathbf{L}$. We compare four algorithms: (i) Q-MC, (ii) Q-MC given the true quantization bin boundaries (denoted Q-MCb), (iii) real-valued MC with $Y_{i,j} \in \{1, \ldots, P\}$ mapped to their corresponding real values as input (denoted MCY), and (iv) real-valued MC using the noisy version of the matrix $\mathbf{L}$ (denoted MCL).

Figure 1 shows the normalized recovery error $E_{\mathbf{L}}$ versus the number of quantization bins P for rank($\mathbf{L}$) = 5 and 10. We can see that Q-MC outperforms real-valued MC in recovering the underlying variable matrix $\mathbf{L}$, especially for coarsely-quantized observations $Y_{i,j}$, i.e., for $P \leq 10$ bins. Furthermore, the performance of Q-MC and Q-MC with known quantization bin boundaries (Q-MCb) is almost equivalent, implying that Q-MC is able to accurately learn the quantization bins. In addition, we see that the performance of Q-MC gradually approaches (with increasing P) the performance of MCL, which operates on the unquantized, real-valued noisy version of the low-rank matrix $\mathbf{L}$.

4.2. Real-world data

We next showcase the performance of Q-MC and Q-RPCA on two real-world collaborative filtering datasets: (i) the MovieLens dataset [24] and (ii) the Dating Agency dataset [14]. The MovieLens dataset is a pruned version consisting of $n = 943$ users rating $m = 1152$ movies; here, only 9% of the ratings are observed. The number of quantized values is $P = 5$ (ratings from 1 to 5). The Dating Agency dataset is pruned to $n = 1000$ users rating $m = 1000$ profiles, with 29.5% ratings observed. The number of quantized values is $P = 10$ (ratings

from 1 to 10). To assess the recovery performance, we use the prediction mean absolute error (MAE) and the prediction root mean square error (RMSE), defined as follows [10]:

$$MAE = \frac{1}{|\overline{\Omega}_{\text{obs}}|} \sum_{(i,j) \in \overline{\Omega}_{\text{obs}}} |Y_{i,j} - \widehat{Y}_{i,j}|,$$

$$RMSE = \sqrt{\sum_{(i,j) \in \overline{\Omega}_{\text{obs}}} (Y_{i,j} - \widehat{Y}_{i,j})^2 / |\overline{\Omega}_{\text{obs}}|}.$$

Here, $\overline{\Omega}_{\text{obs}}$ represents the set of indices corresponding to the unobserved ratings. The predicted entries are obtained by computing $\widehat{Y}_{i,j} = \mathbb{E}[Y_{i,j}|L_{i,j}, \omega_0, \ldots, \omega_P]$. We compare eight different algorithms: Q-MC and Q-RPCA, with both algorithms estimating a set of quantization bin boundaries for (i) the entire dataset (one set of quantization bins), (ii) each item (movie/profile), (iii) each user, and real-valued MC and RPCA. We only use the inverse logit link as it results in lower computational complexity than the inverse probit link.

Table 1 shows the mean of both performance metrics over 25 trials, using 5-fold cross-validation. We can see that Q-MC and Q-RPCA perform similarly, while both methods outperform real-valued MC and RPCA. It is worth noting that the computational complexity of Q-MC and Q-RPCA are close to that of regular MC and RPCA. We further note that the proposed Q-MC and Q-RPCA methods automatically determine the rank of $\mathbf{L}$, which is in contrast to previous *factor analysis*-based methods, which require the rank of $\mathbf{L}$ to be pre-specified (see, e.g., [9–11, 25]).

5. CONCLUSIONS

We have shown that Q-MC and Q-RPCA outperform MC and RPCA on recovering an unknown low-rank matrix $\mathbf{L}$ given quantized and (possibly) sparsely corrupted measurements. Thus, in applications where coarsely quantized (or ordinal) measurements are observed (e.g., with less than about 10 quantization bins), Q-MC and Q-RPCA should be favored over conventional, real-valued MC and RPCA. We emphasize that corresponding theoretical recovery guarantees for Q-MC and Q-RPCA along the lines of [13] are unavailable at this time; this analysis is part of on-going work.

6. REFERENCES

[1] E. J. Candès and T. Tao, "The power of convex relaxation: Near-optimal matrix completion," *IEEE Trans. on Information Theory*, vol. 56, no. 5, pp. 2053–2080, May 2010.

[2] E. J. Candès, X. Li, Y. Ma, and J. Wright, "Robust principal component analysis?," *J. of the ACM*, vol. 58, no. 3, pp. 11, May 2011.

[3] A. E. Waters, A. C. Sankaranarayanan, and R. G. Baraniuk, "SpaRCS: Recovering low-rank and sparse matrices from compressive measurements," in *Advances in Neural Information Processing Systems*, Dec. 2011, pp. 1089–1097.

[4] R. M. Bell and Y. Koren, "Lessons from the Netflix prize challenge," *ACM SIGKDD Explorations Newsletter*, vol. 9, no. 2, pp. 75–79, Dec. 2007.

[5] Y. Koren, R. Bell, and C. Volinsky, "Matrix factorization techniques for recommender systems," *Computer*, vol. 42, no. 8, pp. 30–37, Aug. 2009.

[6] V. Chandrasekaran, S. Sanghavi, P. A. Parrilo, and A. S. Willsky, "Rank-sparsity incoherence for matrix decomposition," *SIAM J. on Optimization*, vol. 21, no. 2, pp. 572–596, June 2011.

[7] D. L. Donoho, "Compressed sensing," *IEEE Trans. on Information Theory*, vol. 52, no. 4, pp. 1289–1306, Apr. 2006.

[8] E. Wang, E. Salazar, D. Dunson, and L. Carin, "Spatio-temporal modeling of legislation and votes," *Bayesian Analysis*, vol. 8, no. 1, pp. 233–268, Mar. 2013.

[9] A. S. Lan, A. E. Waters, C. Studer, and R. G. Baraniuk, "Sparse factor analysis for learning and content analytics," *J. Machine Learning Research*, 2014, to appear.

[10] N. Delannay and M. Verleysen, "Collaborative filtering with interlaced generalized linear models," *Neurocomputing*, vol. 71, no. 7, pp. 1300–1310, Mar. 2008.

[11] Y. Koren and J. Sill, "OrdRec: an ordinal model for predicting personalized item rating distributions," in *Proc. of the 5th ACM Conf. on Recommender Systems*, Oct. 2011, pp. 117–124.

[12] A. S. Lan, C. Studer, A. E. Waters, and R. G. Baraniuk, "Tag-aware ordinal sparse factor analysis for learning and content analytics," in *Proc. 6th Intl. Conf. on Educational Data Mining*, July 2013, pp. 90–97.

[13] M. A. Davenport, Y. Plan, E. van den Berg, and M. Wootters, "1-bit matrix completion," *arXiv preprint:1209.3672v2*, Sep. 2013.

[14] L. Brožovskỳ and V. Petříček, "Recommender system for on-line dating service," in *Proc. Znalosti 2007 Conference*, Feb. 2007.

[15] E. Salazar, M. Cain, E. Darling, S. Mitroff, and L. Carin, "Inferring latent structure from mixed real and categorical relational data," in *Proc. Intl. Conf. on Machine Learning*, Jun. 2012, pp. 1039–1046.

[16] A. Beck and M. Teboulle, "A fast iterative shrinkage-thresholding algorithm for linear inverse problems," *SIAM J. on Imaging Science*, vol. 2, no. 1, pp. 183–202, Mar. 2009.

[17] T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning*, Springer, New York, 2010.

[18] M. Fazel, H. Hindi, and S. P. Boyd, "A rank minimization heuristic with application to minimum order system approximation," in *Proc. American Control Conf.*, June 2001, vol. 6, pp. 4734–4739.

[19] S. Boyd and L. Vandenberghe, *Convex Optimization*, Cambridge University Press, 2004.

[20] J. Cai, E. J. Candès, and Z. Shen, "A singular value thresholding algorithm for matrix completion," *SIAM J. Optimization*, vol. 20, no. 4, pp. 1956–1982, Mar. 2010.

[21] Y. Xu and W. Yin, "A block coordinate descent method for multi-convex optimization with applications to nonnegative tensor factorization and completion," Tech. Rep., Rice University CAAM, Sep. 2012.

[22] J. Nocedal and S. J. Wright, *Numerical Optimization*, Springer Verlag, 1999.

[23] J. Duchi, S. Shalev-Shwartz, Y. Singer, and T. Chandra, "Efficient projections onto the $\ell 1$-ball for learning in high dimensions," in *Proc. 25th Intl. Conf. on Machine Learning*, July 2008, pp. 272–279.

[24] J. L. Herlocker, J. A. Konstan, A. Borchers, and J. Riedl, "An algorithmic framework for performing collaborative filtering," in *Proc. 22nd Annual Intl. ACM SIGIR Conf. on Research and Development in Information Retrieval*, Aug. 1999, pp. 230–237.

[25] A. Zymnis, S. Boyd, and E. Candès, "Compressed sensing with quantized measurements," *IEEE Sig. Proc. Letters*, vol. 17, no. 2, pp. 149–152, Feb. 2010.