ONE-BIT MATRIX COMPLETION FOR PAIRWISE COMPARISON MATRICES

ANDREW K. MASSIMINO AND MARK A. DAVENPORT*

Abstract. In this paper we consider the related problems of ranking and of recovering a matrix of pairwise comparisons from binary observations. We describe a naïve adaptation of the one-bit matrix completion framework, but then note that additional constraints that arise in the context of ranking allows us to replace nuclear norm minimization with a more direct approach. This ultimately leads to a novel viewpoint on a classic approach to the ranking problem. Both theoretical and experimental results show that this simplified approach to recovering a pairwise comparison matrix performs significantly better than the naïve approach.

Key words. matrix completion, ranking, pairwise comparisons, maximum likelihood, Bradley-Terry-Luce choice model

1. Introduction. The problem of ranking a set of items based only on comparisons between pairs of items arises in a wide range of applications. The items could be consumer products, sports teams, tweets or online forum comments, candidate drug therapies, job applications, academic papers or proposals, or any other collection of objects or courses of action. In this paper we investigate the connection between ranking and the problem of recovering a low-rank matrix from an incomplete set of observations, or *matrix completion* [1]. To see the connection between these two problems, assume that each of the n items to be ranked can be assigned a numerical rating (or score) $r_i \in \mathbb{R}$ and that item i is preferred to item j if $r_i > r_j$. These ratings taken together form the rating vector $\boldsymbol{r} \in \mathbb{R}^n$. Next consider the matrix $\boldsymbol{M} \in \mathbb{R}^{n \times n}$ in which each element M_{ij} represents the difference between the ratings for items i and j, i.e., $M_{ij} = r_i - r_j$. Equivalently, we can write $\boldsymbol{M} = \boldsymbol{r}\boldsymbol{e}^T - \boldsymbol{e}\boldsymbol{r}^T$, where $\boldsymbol{e} = [1, 1, \ldots, 1]^T$.

Clearly, given the matrix $\boldsymbol{M}$ we could directly recover the vector $\boldsymbol{r}$, giving us a rank-ordered list of all items. However, in practice we rarely have access to the matrix $\boldsymbol{M}$. Obtaining this matrix would require comparisons between every possible pair of items, which can be prohibitive in practice for even moderate sized n. For example, in constructing a ranking of sports teams, we can rarely expect every team to play every other possible team. Moreover, in most settings the outcomes of these comparisons will be extremely noisy and are often highly "quantized"—we may only observe which team wins, and have no strict assurance that the winning team was in fact the superior one. Similar problems arise whenever the comparisons are the result of human judgements, which can be notoriously unreliable and may only tell us which of two items was preferred. In this paper we consider the question of whether the underlying matrix $\boldsymbol{M}$ can be recovered from such noisy one-bit observations when the observations are potentially highly incomplete. We approach the problem from the perspective of *one-bit matrix completion*, which asserts that since the matrix $\boldsymbol{M}$ is low-rank, we should be able to recover $\boldsymbol{M}$ accurately via a simple convex program [2].

In the remainder of this paper, we will first review the theory of one-bit matrix completion and describe its implications in the context of ranking. By imposing additional constraints, we obtain improved theoretical results for an algorithm that turns out to be essentially equivalent to a classical approach to ranking. We conclude with a set of experiments that illustrate the performance of our proposed approach.

*Georgia Institute of Technology, School of Electrical and Computer Engineering, Atlanta, GA. Email: {massimino,mdav}@gatech.edu

2. One-Bit Matrix Completion. We begin with a brief overview of the one-bit matrix completion framework. Suppose that $\boldsymbol{Y}$ denotes a complete set of noisy, binary comparisons given by

$$Y_{ij} = \begin{cases} +1 \text{ (item } i \text{ preferred)} & \text{with probability } f(M_{ij}) \\ -1 \text{ (item } j \text{ preferred)} & \text{with probability } 1 - f(M_{ij}) \end{cases} \tag{2.1}$$

where $f(M_{i,j}) = P(M_{ij} > 0)$. In our context, this corresponds to the probability that item i is preferred to item j, and the choice of f would determine the chance that the outcomes are "upsets" with respect to the underlying ratings. In this paper we will focus on the case where f is given by the standard logistic function $f(x) = (1+e^{-x})^{-1}$, but other natural choices are also possible (see [2]). In the one-bit matrix completion setting, we assume that we are able to observe $\boldsymbol{Y}$ on a subset of indices indexed by Ω. In this case, if we let Ω^+ denote the subset of Ω where $Y_{ij} = +1$ and similarly for Ω^-, then the log-likelihood function is given by

$$F_\Omega(\boldsymbol{X}) = \sum_{(i,j)\in\Omega^+} \log f(X_{ij}) + \sum_{(i,j)\in\Omega^-} \log f(1 - X_{ij}). \tag{2.2}$$

If the underlying matrix $\boldsymbol{M}$ has rank s and $\|\boldsymbol{M}\|_\infty \le \alpha$, then the approach advocated in [2] is to set

$$\widehat{\boldsymbol{M}} = \arg\max_{\boldsymbol{X}} F_\Omega(\boldsymbol{X}) \quad \text{subject to} \quad \|\boldsymbol{X}\|_* \le \alpha n\sqrt{s} \text{ and } \|\boldsymbol{X}\|_\infty \le \alpha. \tag{2.3}$$

In the case where $\boldsymbol{M} = \boldsymbol{r}\boldsymbol{e}^T - \boldsymbol{e}\boldsymbol{r}^T$, we have that $s = 2$, and hence we obtain the following bound on $\|\boldsymbol{M} - \widehat{\boldsymbol{M}}\|_F$ as a special case of Theorem 1 in [2].

Theorem 2.1. *Let $f(x) = (1 + e^{-x})^{-1}$ and suppose that we obtain observations of the form in (2.1) for indices $(i,j) \in \Omega$, where Ω is a set of $m \ge 4n\log n$ elements chosen uniformly at random and $\boldsymbol{M} = \boldsymbol{r}\boldsymbol{e}^T - \boldsymbol{e}\boldsymbol{r}^T$ with $r_i \in [0, \rho]$ for $i = 1, \dots, n$. Then the solution to (2.3) with $\alpha = \rho$, $s = 2$ will satisfy*

$$\frac{1}{n^2}\|\boldsymbol{M} - \widehat{\boldsymbol{M}}\|_F^2 \le C_1\rho\sqrt{\frac{n}{m}},$$

with probability at least $1 - C_2/n$, where C_1 and C_2 are absolute constants.

Note that the requirement that $m = O(n\log n)$ is not a particularly strong assumption since as a consequence of the *Coupon collector problem* we need at least $n\log n$ random observations just to ensure that we obtain at least one comparison per item in our set. Next observe that from the estimate $\widehat{\boldsymbol{M}}$, the rating vector $\boldsymbol{r}$ can be easily estimated (up to an order-preserving affine shift) via

$$\boldsymbol{r} = \frac{(\widehat{\boldsymbol{M}} - \widehat{\boldsymbol{M}}^T)\boldsymbol{e}}{2n}. \tag{2.4}$$

This essentially corresponds to the well-known *Borda count*, and it is not difficult to show (e.g., see [4, 5]) that the $\boldsymbol{r}$ defined in (2.4) satisfies

$$\boldsymbol{r} = \arg\min_{\boldsymbol{r}} \left\| \left(\boldsymbol{r}\boldsymbol{e}^T - \boldsymbol{e}\boldsymbol{r}^T\right) - \widehat{\boldsymbol{M}} \right\|_F.$$

This is certainly one viable approach. However, this naïve application of generic matrix completion techniques ignores the fact that in this specific context, the matrix

$\boldsymbol{M}$ has considerable additional structure beyond having low rank, and it may be possible to leverage this structure to obtain improved results. For example, in the context of rank aggregation (but without quantized measurements), [4] considers a method that incorporates the requirement that $\boldsymbol{M}$ is skew-symmetric directly as a constraint in the recovery algorithm. This seems more natural than recovering a generic $\widehat{\boldsymbol{M}}$ and then after the fact estimating a skew-symmetric approximation to $\widehat{\boldsymbol{M}}$ as in the approach described above. Inspired by [4], we will now consider an alternative approach to adapting the one-bit matrix completion framework to the ranking problem.

3. An Alternative Approach. We begin by noting that when $\boldsymbol{X} = \boldsymbol{x}\boldsymbol{e}^T - \boldsymbol{e}\boldsymbol{x}^T$ we can rewrite the log likelihood function in (2.2) as

$$F_\Omega(\boldsymbol{x}) = \sum_{(i,j)\in\Omega^+} \log f(x_i - x_j) + \sum_{(i,j)\in\Omega^-} \log f(1 - (x_i - x_j)). \tag{3.1}$$

Next we note that in the case of the logistic model (or any other symmetric distribution) we have $\log f(x_i - x_j) = \log(1 - f(x_j - x_i))$. Thus, since whenever $(i,j) \in \Omega^+$ we automatically have $(j,i) \in \Omega^-$, in this case we can reduce (3.1) to simply $F_\Omega(\boldsymbol{x}) = 2F_{\Omega^+}(\boldsymbol{x})$, where

$$F_{\Omega^+}(\boldsymbol{x}) = \sum_{(i,j)\in\Omega^+} \log f(x_i - x_j). \tag{3.2}$$

Finally, if we allow Ω to be a multiset where a pair (i,j) can appear multiple times (i.e., we allow for repeated comparisons between a pair of objects), then we can define $A_{ij} = |\{(i,j) \in \Omega^+\}|$ and consider

$$F_{\Omega^+}(\boldsymbol{x}) = \sum_{(i,j)} A_{ij} \log f(x_i - x_j). \tag{3.3}$$

We now consider the solution to the following optimization problem:

$$\widehat{\boldsymbol{r}} = \arg\max_{\boldsymbol{x}} F_{\Omega^+}(\boldsymbol{x}) \quad \text{subject to} \quad x_i \geq 0 \text{ for } i = 1, \ldots, n \text{ and } \|\boldsymbol{x}\|_1 \leq \rho. \tag{3.4}$$

We can obtain $\widehat{\boldsymbol{M}}$ from $\widehat{\boldsymbol{r}}$ via $\widehat{\boldsymbol{M}} = \widehat{\boldsymbol{r}}\boldsymbol{e}^T - \boldsymbol{e}\widehat{\boldsymbol{r}}^T$. The solution to this optimization problem will satisfy the following guarantee.

THEOREM 3.1. *Let $f(x) = (1 + e^{-x})^{-1}$ and suppose that we obtain observations of the form in (2.1) for indices $(i,j) \in \Omega$, where Ω is a set of $m \geq 2n \log n$ pairs of items chosen uniformly at random and $\boldsymbol{M} = \boldsymbol{r}\boldsymbol{e}^T - \boldsymbol{e}\boldsymbol{r}^T$ with $r_i \geq 0$ for $i = 1, \ldots, n$ and $\|\boldsymbol{r}\|_1 \leq \rho$. Then with probability at least $1 - C_1/n$, the solution to (3.4) with $\widehat{\boldsymbol{M}} = \widehat{\boldsymbol{r}}\boldsymbol{e}^T - \boldsymbol{e}\widehat{\boldsymbol{r}}^T$ will satisfy*

$$\frac{1}{\binom{n}{2}} \|\boldsymbol{M} - \widehat{\boldsymbol{M}}\|_F^2 \leq C_2 \rho \sqrt{\frac{1}{m}},$$

where C_1 and C_2 are absolute constants.

Proof. The proof follows the same structure as the proof of Theorem 1 in [2], so we provide only an outline of the proof here. The proof begins by showing that there exists a constant C such that

$$\|\boldsymbol{M} - \widehat{\boldsymbol{M}}\|_F^2 \leq C \sum_{i,j} D(f(M_{ij}) \| f(\widehat{M}_{ij})), \tag{3.5}$$

where $D(\cdot\|\cdot)$ denotes the Kullback-Leibler divergence. Next we define

$$G = \{\boldsymbol{x} \in \mathbb{R}^n : x_i \geq 0 \text{ for } i = 1, \ldots, n \quad \text{and} \quad \|\boldsymbol{x}\|_1 \leq \rho\}.$$

Using the same argument as in [2] one can show that if we define $H(\boldsymbol{x}) = F_{\Omega^+}(\boldsymbol{x}) - \mathbb{E}F_{\Omega^+}(\boldsymbol{x})$, where the expectation is over both the choice of Ω and the outcomes $\boldsymbol{Y}$, then as a consequence of the fact that $F_{\Omega^+}(\widehat{\boldsymbol{r}}) \geq F_{\Omega^+}(\boldsymbol{r})$ we have

$$\frac{1}{\binom{n}{2}} \sum_{i,j} D(f(M_{ij}) \| f(\widehat{M}_{ij})) \leq \frac{2}{m} \sup_{\boldsymbol{x} \in G} |H(\boldsymbol{x})|.$$

Combining this with (3.5) and Lemma 3.2 establishes the result with $C_2 = C \cdot C_0$. ∎

LEMMA 3.2. *There exist constants C_0 and C_1 such that*

$$P\left(\sup_{\boldsymbol{x} \in G} |H(\boldsymbol{x})| \geq C_0 \rho \sqrt{m + n \log n}\right) \leq \frac{C_1}{n}.$$

Proof. Again, the proof follows the same structure as the proof of Lemma 1 in [2], so we provide only a rough outline. We begin by noting that as a consequence of Markov's inequality we have that for $h > 0$

$$P\left(\sup_{\boldsymbol{x} \in G} |H(\boldsymbol{x})| \geq C_0 \rho \sqrt{m + n \log n}\right) \leq \frac{\mathbb{E}\left[\sup_{\boldsymbol{x} \in G} |H(\boldsymbol{x})|^h\right]}{\left(C_0 \rho \sqrt{m + n \log n}\right)^h} \tag{3.6}$$

Following the same line of reasoning as in [2], one can show that

$$\mathbb{E}\left[\sup_{\boldsymbol{x} \in G} |H(\boldsymbol{x})|^h\right] \leq 4^h \|\boldsymbol{x}\boldsymbol{e}^T - \boldsymbol{e}\boldsymbol{x}^T\|_*^h \mathbb{E}\left[\|\boldsymbol{E} \circ \boldsymbol{A}\|^h\right], \tag{3.7}$$

where $\boldsymbol{E}$ is a matrix with i.i.d. Rademacher random variables. Using the same arguments as in [2] we can show that there exists a constant c such that

$$\mathbb{E}\left[\|\boldsymbol{E} \circ \boldsymbol{A}\|^h\right] \leq \left(c\sqrt{\frac{m + n \log n}{n}}\right)^h. \tag{3.8}$$

We also have that

$$\|\boldsymbol{x}\boldsymbol{e}^T - \boldsymbol{e}\boldsymbol{x}^T\|_* \leq \|\boldsymbol{x}\boldsymbol{e}^T\|_* + \|\boldsymbol{e}\boldsymbol{x}^T\|_* = 2\sqrt{n}\|\boldsymbol{x}\|_2 \leq 2\sqrt{n}\rho. \tag{3.9}$$

Combining (3.7), (3.8), and (3.9), we obtain

$$\mathbb{E}\left[\sup_{\boldsymbol{x} \in G} |H(\boldsymbol{x})|^h\right] \leq \left(8c\rho\sqrt{m + n \log n}\right)^h.$$

Plugging this in to (3.6) and setting $h = \log n$ establishes the lemma. ∎

4. Connection to Previous Approaches to Ranking. Using A_{ij} as previously defined, we can now write the log likelihood as

$$F_{\Omega^+}(\boldsymbol{r}) = \sum_{i<j} A_{ij} \log(f(r_i - r_j)) + A_{ji} \log(f(r_j - r_i)). \tag{4.1}$$

In the case of the logistic function, we have

$$\begin{aligned} F_{\Omega^+}(\boldsymbol{r}) &= -\sum_{i<j} A_{ij}\log(e^{(r_j-r_i)}+1) + A_{ji}\log(e^{(r_i-r_j)}+1) \\ &= \sum_{i<j} \log\left[(e^{r_j}/e^{r_i}+1)^{-A_{ij}}\,(e^{r_i}/e^{r_j}+1)^{-A_{ji}}\right]. \end{aligned} \tag{4.2}$$

Up to a set of monotonic transformations in $\boldsymbol{r}$, our model is actually equivalent to the classic *Bradley-Terry-Luce* model, which has been studied extensively in the context of ranking:

$$F_{\Omega^+}(\boldsymbol{w} = e^{\boldsymbol{r}}) = \log \prod_{i<j} \left(\frac{w_i}{w_i + w_j}\right)^{A_{ij}} \left(\frac{w_j}{w_j + w_i}\right)^{A_{ji}}. \tag{4.3}$$

From this perspective, one can view our proposed algorithm as only a slight variant of a classic approach described in [3], with the difference being in the particular form of our constraints on $\boldsymbol{r}$. Specifically, both our approach as well as the approach described in [3] can be expressed as

$$\text{maximize} \quad F_{\Omega^+}(\boldsymbol{w}) \quad \text{subject to} \quad h(\boldsymbol{w}) = 0. \tag{4.4}$$

The choice of constraints is somewhat arbitrary, but sensible examples include $h(\boldsymbol{w}) = 1 - \sum_{i=1}^n w_i$ for $w_i > 0$ as in [3] and $h(e^{\boldsymbol{r}}) = 1 - \sum_{i=1}^n r_i$ for $r_i \geq 0$ as in Section 3 .

5. Synthetic Simulation. To evaluate our approach, we perform a set of numerical simulations that compares the performance of our proposed approach with the naïve adaptation of one-bit matrix completion on a synthetic example where the true underlying ratings are known. In this experiment, we generate an evenly spaced ranking vector $\boldsymbol{r} \in \mathbb{R}^n$ such that $\|\widehat{\boldsymbol{r}}\|_1 = \sum_i r_i = \rho$. We observe random match-ups between items and these binary measurements are subjected to noise according to the logistic model. Both the number of measurements and the measurement noise are varied. The estimated ranking vector $\widehat{\boldsymbol{r}}$ is computed using the `SPGL1` [6] solver:

To compare the fidelity of reconstruction, we use a distance $K_d \colon \mathbb{R}^n \times \mathbb{R}^n \to [0,1]$ which counts the number of discordant pairs (discrepancies between ranking lists) relative to the total number of pairs of items. Discordant pairs can be thought of thought of as how many "swaps" it would take to transform one ranking list to another.

$$K_d(r, \widehat{r}) = \frac{|\{(r_i > r_j \text{ and } \widehat{r}_i < \widehat{r}_j) \text{ or } (r_i < r_j \text{ and } \widehat{r}_i > \widehat{r}_j)\}|}{n(n-1)/2} \tag{5.1}$$

This function is related to the *Kendall-τ*, which is a measure of how close two ranking vectors are [5]. As we have defined it, 0 represents complete agreement and 1 is full disagreement. Other metrics such as mean squared error were considered, but in ranking, the order of the sorted ranking list is typically more important than the precise values of the ranking vector. Results of this experiment are given in Fig. 5.1, which shows the log of the Kendall-τ distance for varying noise and number of observations (as a percentage). Increasing $\sum_i r_i$ implies *decreasing* noise as is consistent with the asymptotically decreasing trend in the τ-distance. The results show that the modification made in Section 3 significantly improves the recovery of ranking vectors.

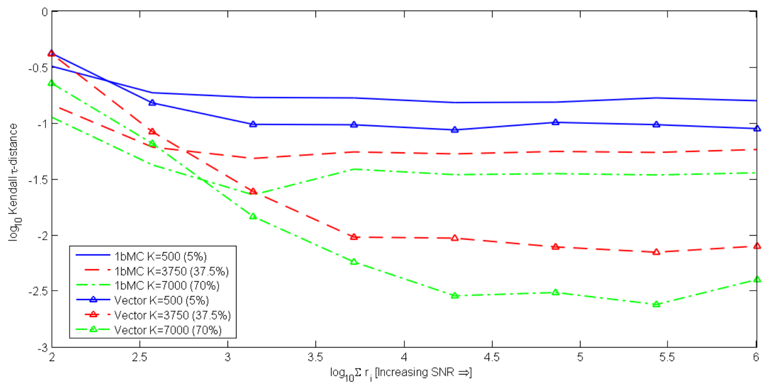

FIG. 5.1. *Comparison of synthetic experiment results for one-bit matrix completion and the simplified model (labeled vector).*

6. Conclusion. In this paper, we have shown that ranking from pairwise comparisons can be viewed as a matrix completion problem. However, simplification due to the matrix structure in this context allows us to dispense with nuclear norm minimization and reduce the problem to a classical approach to ranking. We thus provide new theoretical insight into this traditional approach to ranking. We find that this proposed approach leads to significant recovery performance gains as compared to the aforementioned matrix completion solution. Specifically, there is an improvement by a factor of $\sqrt{n}$ in the theoretical mean square recovery error between the two approaches. This is supported by our experimental results.

In this paper, we considered only skew-symmetric, rank 2 matrices of the form $\boldsymbol{M} = \boldsymbol{r}\boldsymbol{e}^T - \boldsymbol{e}\boldsymbol{r}^T$. Although this simplifies the problem greatly, it is likely that many inferences that can be made with the data are lost in imposing such a constraint. For example, in general there may be multiple factors that determine which item is preferred, and so in some cases it may be desirable to consider matrices M of higher rank. Thus, an important area of future work in applying matrix completion to pairwise ranking may involve loosening the assumptions made on M.

REFERENCES

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

[2] M. Davenport, Yaniv Plan, Ewout van den Berg, and Mary Wootters, *1-bit matrix completion*, Arxiv preprint arxiv:1209.3672, (2012).

[3] L. Ford, *Solution of a ranking problem from binary comparisons*, Amer. Math. Monthly, 64 (1957), pp. 28–33.

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

[5] A. Langville, *Who's #1?: The science of rating and ranking*, Princeton University Press, Princeton N.J, 2012.

[6] E. van den Berg and M. P. Friedlander, *SPGL1: A solver for large-scale sparse reconstruction*, June 2007. http://www.cs.ubc.ca/labs/scl/spgl1.