Supplemental Material

A. Generalization of the ASR algorithm with Regularization

In this section, we shall present a generalized version of the ASR algorithm that relaxes the assumption that each set S_a is of the same fixed cardinality m, and each set S_a is compared the same number of times L. The intuition behind this generalization is that each comparison carries an equal amount of information, and thus, we should give a higher preference to the empirical estimates $\hat{p}_{i|S_a}$ corresponding to sets with more comparisons. Furthermore, comparisons on smaller sets are more reliable than comparisons on larger sets. In general, sets with larger cardinality should have proportionately more comparisons. Lastly, in practice, we often encounter comparison data for which the random walk $\hat{\mathbf{P}}$ on the comparison graph G_c is not strongly connected. We can resolve this issue through regularization. With these in mind, we update our algorithm as discussed below:

Given general comparison data $\mathbf{Y}' = \{(S_a, \mathbf{y}_a)_{a=1}^d\}$, where $S_a \subseteq [n]$ is of cardinality $|S_a|$, and $\mathbf{y}_a = (y_a^1, \ldots, y_a^{L_a})$, we define d'_i for each $i \in [n]$ as

$$d'_i := \sum_{a\in[d]:i\in S_a} \left(\frac{L_a}{|S_a|} + \lambda \right)$$

where λ is a regularization parameter. Intuitively, one can think of the regularization as adding $\lambda|S_a|$ pseudo-comparisons to each set S_a, with each item in the set winning an equal λ times. Furthermore, we define $n_{i|S_a}$ to be the number of times item $i \in S_a$ won in a $|S_a|$-way comparison amongst items in S_a, i.e. for all $a \in [d]$, for all $i \in S_a$,

$$n_{i|S_a} := \sum_{l=1}^{L_a} \mathbf{1}[y_a^l = i] \tag{7}$$

Using the above notation, we set up a Markov chain $\hat{\mathbf{P}}' \in \mathbb{R}_+^{n\times n}$ such that entry (i, j) is

$$\hat{P}'_{ij} := \frac{1}{d'_i} \sum_{a\in[d]:i,j\in S_a} \left(\frac{n_{j|S_a} + \lambda}{|S_a|} \right) \tag{8}$$

One can verify that this non-negative matrix is indeed row stochastic, hence corresponds to the transition matrix of a Markov chain. One can also verify that this construction reduces to a regularized version of $\hat{\mathbf{P}}$ (Eq. (2)) when all sets are of an equal size and are compared an equal number of times, and is identical to $\hat{\mathbf{P}}$ when $\lambda = 0$. Lastly, we define the matrix $\mathbf{D}'$ as a diagonal matrix, with diagonal entry $D'_{ii} := d'_i$, $\forall i \in [n]$. Similar to ASR, we compute the stationary distribution of $\hat{\mathbf{P}}'$, and output a (normalized) $\mathbf{D}'^{-1}$ transform of this stationary distribution.

Algorithm 3 Generalized-ASR

Input Markov chain $\hat{\mathbf{P}}'$ (according to Eq. (8))
Initialize $\hat{\boldsymbol{\pi}} = (\frac{1}{n}, \cdots, \frac{1}{n})^\top \in \Delta_n$
while estimates do not converge **do**
  $\hat{\boldsymbol{\pi}}' \leftarrow \hat{\mathbf{P}}'^\top \hat{\boldsymbol{\pi}}'$
end while
Output $\hat{\mathbf{w}}' = \frac{\mathbf{D}'^{-1}\hat{\boldsymbol{\pi}}'}{\|\mathbf{D}'^{-1}\hat{\boldsymbol{\pi}}'\|_1}$

B. Proof of Proposition 1

Proposition 1. *Given items $[n]$ and comparison data $\mathbf{Y} = \{(S_a, \mathbf{y}_a)\}_{a=1}^d$, let $\hat{\boldsymbol{\pi}}$ be the stationary distribution of the Markov chain $\hat{\mathbf{P}}$ constructed by* ASR*, and let $\hat{\mathbf{w}}^{LSR}$ be the stationary distribution of the Markov chain $\hat{\mathbf{P}}^{LSR}$. Then $\hat{\mathbf{w}}^{LSR} = \frac{\mathbf{D}^{-1}\hat{\boldsymbol{\pi}}}{\|\mathbf{D}^{-1}\hat{\boldsymbol{\pi}}\|_1}$. The same result is also true for $\hat{\mathbf{w}}^{RC}$ for the case of pairwise comparisons.*

Proof. Consider the estimates $\hat{\mathbf{w}} = \mathbf{D}^{-1}\hat{\boldsymbol{\pi}}/\|\mathbf{D}^{-1}\hat{\boldsymbol{\pi}}\|_1$ returned by the ASR algorithm upon convergence. In order to prove this lemma it is sufficient to prove that $\mathbf{D}\hat{\mathbf{w}}^{\text{LSR}}$ is an invariant measure (an eigenvector associated with eigenvalue 1) of the Markov chain $\hat{\mathbf{P}}$ corresponding to the ASR algorithm.

Since $\hat{\mathbf{w}}^{\text{LSR}}$ is the stationary distribution (also an eigenvector corresponding to eigenvalue 1) of $\hat{\mathbf{P}}^{\text{LSR}}$, we have

$$\hat{\mathbf{w}}^{\text{LSR}} = (\hat{\mathbf{P}}^{\text{LSR}})^\top \hat{\mathbf{w}}^{\text{LSR}}.$$

Following the definition (Eq. (4)) of $\hat{\mathbf{P}}^{\text{LSR}}$, we have the following relation for all $1 \le i \le n$

$$\begin{aligned}
\hat{w}_i^{\text{LSR}} &= \hat{w}_i^{\text{LSR}} \left(1 - \epsilon \sum_{j\neq i} \sum_{a:i,j\in S_a} p_{j|S_a} \right) \\
&\quad + \epsilon \sum_{j\neq i} \sum_{a:i,j\in S_a} p_{j|S_a} \hat{w}_j^{\text{LSR}} \\
\implies \sum_{j\neq i} \sum_{a:i,j\in S_a} p_{j|S_a} \hat{w}_i^{\text{LSR}} &= \sum_{j\neq i} \sum_{a:i,j\in S_a} p_{j|S_a} \hat{w}_j^{\text{LSR}}.
\end{aligned}$$

We shall use this relation to prove that $\hat{\mathbf{P}}^\top \mathbf{D}\hat{\mathbf{w}}^{\text{LSR}} = \mathbf{D}\hat{\mathbf{w}}^{\text{LSR}}$, where $\hat{\mathbf{P}}$ is the transition matrix corresponding to the Markov chain constructed by ASR. Consider the i^{th}

coordinate $[\hat{\mathbf{P}}^\top \mathbf{D}\hat{\mathbf{w}}^{\text{LSR}}]_i$ of the vector $\hat{\mathbf{P}}^\top \mathbf{D}\hat{\mathbf{w}}^{\text{LSR}}$

$$\begin{aligned}
[\hat{\mathbf{P}}^\top \mathbf{D}\hat{\mathbf{w}}^{\text{LSR}}]_i &= \frac{1}{d_i}\sum_{a:i\in S_a} p_{i|S_a} d_i \hat{w}_i^{\text{LSR}} \\
&\quad + \sum_{j\neq i}\frac{1}{d_j}\sum_{b:i,j\in S_b} p_{j|S_b} d_j \hat{w}_j^{\text{LSR}} \\
&= \sum_{a:i\in S_a} p_{i|S_a}\hat{w}_i^{\text{LSR}} + \sum_{j\neq i}\sum_{b:i,j\in S_b} p_{j|S_b}\hat{w}_i^{\text{LSR}} \\
&= \sum_{a:i\in S_a}\Big(\sum_{j\in S_a} p_{j|S_a}\Big)\hat{w}_i^{\text{LSR}} \\
&= \sum_{a:i\in S_a}\hat{w}_i^{\text{LSR}} \\
&= d_i\hat{w}_i^{\text{LSR}} = [\mathbf{D}\hat{\mathbf{w}}^{\text{LSR}}]_i\,,
\end{aligned}$$

where the second equality follows from the relation we proved earlier. Furthermore, this identity holds for all $1 \le i \le n$, from which we can conclude $\hat{\mathbf{P}}^\top \mathbf{D}\hat{\mathbf{w}}^{\text{LSR}} = \mathbf{D}\hat{\mathbf{w}}^{\text{LSR}}$. Furthermore, if the respective Markov chains induced by the comparison data are ergodic, then the corresponding stationary distributions must be unique, which is sufficient to prove both LSR and ASR return the same estimates upon convergence.

Since Luce spectral ranking is a generalization of the rank centrality algorithm, the transition matrix $\hat{\mathbf{P}}^{\text{LSR}}$ is identical to the transition matrix $\hat{\mathbf{P}}^{\text{RC}}$ in the pairwise comparison setting after setting $\epsilon = \frac{1}{d_{\max}}$, and thus, we can also conclude $\hat{\mathbf{P}}^\top \mathbf{D}\hat{\mathbf{w}}^{\text{RC}} = \mathbf{D}\hat{\mathbf{w}}^{\text{RC}}$. Thus, the statement of the lemma follows. □

C. Proof of Proposition 2

Proposition 2. *Let the probability transition matrix $\mathbf{P}$ for our random walk be as defined in Eq. (1). Let $\mathbf{P}^{RC}$ and $\mathbf{P}^{LSR}$ be as defined in Eq. (3) and Eq. (4), respectively. Then*

$$\frac{d_{\min}}{d_{\max}}\mu(\mathbf{P}) \le \mu(\mathbf{P}^{RC}) \le \mu(\mathbf{P})\,,$$

and

$$\epsilon d_{\min}\mu(\mathbf{P}) \le \mu(\mathbf{P}^{LSR}) \le \mu(\mathbf{P})\,,$$

where $\epsilon = O(\frac{1}{d_{\max}})$.

In order to prove this lemma, we will use the following result due to (Diaconis & Saloff-Coste, 1993) which compares the spectral gaps of two reversible random walks.

Lemma 3. *(Diaconis & Saloff-Coste, 1993) Let $\mathbf{Q}$ and $\mathbf{P}$ be reversible Markov chains on a finite set $[n]$ representing random walks on a graph $G = ([n], E)$, i.e. $P_{ij} = Q_{ij} = 0$ for all $(i,j) \notin E$. Let $\boldsymbol{\nu}$ and $\boldsymbol{\pi}$ be the stationary distributions of $\mathbf{Q}$ and $\mathbf{P}$, respectively. Then the spectral gaps of $\mathbf{Q}$ and $\mathbf{P}$ are related as*

$$\frac{\mu(\mathbf{P})}{\mu(\mathbf{Q})} \ge \frac{\alpha}{\beta}$$

where $\alpha := \min_{(i,j)\in E}\{\pi_i P_{ij}/\nu_i Q_{ij}\}$ and $\beta := \max_{i\in[n]}\{\pi_i/\nu_i\}$.

We are now ready to prove Proposition 2.

Proof. (of Proposition 2) To prove this lemma, we shall leverage the above comparison lemma due to (Diaconis & Saloff-Coste, 1993), that compares the spectral gaps of two arbitrary reversible Markov Chains. Let $\mathbf{P}$ (Eq. (2)) be the reversible Markov chain corresponding to ASR with stationary distribution $\boldsymbol{\pi} = \mathbf{D}\mathbf{w}/\|\mathbf{D}\mathbf{w}\|_1$, and let $\mathbf{P}^{\text{LSR}}$ (Eq. (4)) be the reversible Markov chain corresponding to LSR (RC in the pairwise case) with stationary distribution $\boldsymbol{\pi}^{\text{LSR}}$. Then by Lemma 3,

$$\frac{\mu(\mathbf{P}^{\text{LSR}})}{\mu(\mathbf{P})} \ge \frac{\alpha}{\beta}$$

where

$$\begin{aligned}
\alpha &:= \min_{(i,j):\exists a \text{ s.t. } i,j\in S_a}\left(\frac{\pi_i^{\text{LSR}}P_{ij}^{\text{LSR}}}{\pi_i P_{ij}}\right), \\
\beta &:= \max_{i\in[n]}\left(\frac{\pi_i^{\text{LSR}}}{\pi_i}\right).
\end{aligned}$$

From the definition of $\mathbf{P}$, and $\mathbf{P}^{\text{LSR}}$, we have

$$\begin{aligned}
\mathbf{P}_{ij} &= \frac{1}{d_i}\sum_{a\in[d]:i,j\in S_a}\frac{w_j}{\sum_{k\in S_a} w_k}, \\
\mathbf{P}_{ij}^{\text{LSR}} &= \epsilon\sum_{a\in[d]:i,j\in S_a}\frac{w_j}{\sum_{k\in S_a} w_k}
\end{aligned}$$

From the above equations and Proposition 1, it is easy to see that

$$\begin{aligned}
\alpha &= \epsilon\|\mathbf{D}\mathbf{w}\|_1, \quad \text{and} \\
\beta &= \frac{\|\mathbf{D}\mathbf{w}\|_1}{d_{\min}} \\
\implies \mu(\mathbf{P}^{\text{LSR}}) &\ge \epsilon d_{\min}(\mu(\mathbf{P}))
\end{aligned}$$

Following an identical line of reasoning, we have

$$\frac{\mu(\mathbf{P})}{\mu(\mathbf{P}^{\text{LSR}})} \ge \frac{\alpha'}{\beta'}$$

where

$$\begin{aligned}
\alpha' &= \min_{(i,j):\exists a \text{ s.t. } i,j\in S_a}\left(\frac{\pi_i P_{ij}}{\pi_i^{\text{LSR}}P_{ij}^{\text{LSR}}}\right), \\
\beta' &= \max_{i\in[n]}\left(\frac{\pi_i}{\pi_i^{\text{LSR}}}\right)
\end{aligned}$$

From the definition of $\mathbf{P}$, and $\mathbf{P}^{\text{LSR}}$, we have

$$\alpha' = \frac{1}{\|\mathbf{D}\mathbf{w}\|_1 \epsilon}, \quad \text{and}$$
$$\beta' = \frac{d_{\max}}{\|\mathbf{D}\mathbf{w}\|_1}$$
$$\implies \mu(\mathbf{P}) \geq \frac{1}{\epsilon d_{\max}}(\mu(\mathbf{P}^{\text{LSR}}))\,.$$

Since $\epsilon \leq 1/d_{\max}$, we get the following comparison between the spectral gaps of the Markov chains corresponding to the two approaches

$$\epsilon d_{\min}\mu(\mathbf{P}) \leq \mu(\mathbf{P}^{\text{LSR}}) \leq \mu(\mathbf{P})\,.$$

The same analysis works for the Markov chain $\mathbf{P}^{\text{RC}}$ constructed by rank centrality for the pairwise comparison case with $\epsilon = 1/d_{\max}$, from which we can conclude

$$\frac{d_{\min}}{d_{\max}}\mu(\mathbf{P}) \leq \mu(\mathbf{P}^{\text{RC}}) \leq \mu(\mathbf{P})\,.$$

□

D. Proof of Theorem 1

Theorem 1. *Given items $[n]$ and comparison data $\mathbf{Y} = \{(S_a, \mathbf{y}_a)\}_{a=1}^d$, let each set S_a of cardinality m be compared L times, with outcomes $\mathbf{y}_a = (y_a^1, \cdots, y_a^L)$ produced as per a MNL model with parameters $\mathbf{w} = (w_1, \ldots, w_n)$, such that $\|\mathbf{w}\|_1 = 1$. If the random walk $\hat{\mathbf{P}}$ (Eq. (2)) on the comparison graph $G_c([n], E)$ induced by the comparison data $\mathbf{Y}$ is strongly connected, then the* ASR *algorithm (Algorithm 1) converges to a unique distribution $\hat{\mathbf{w}}$, which with probability $\geq 1 - 3n^{-(C^2-50)/25}$ satisfies the following error bound*

$$\|\mathbf{w} - \hat{\mathbf{w}}\|_{TV} \leq \frac{C\,\kappa\,d_{\text{avg}}}{\mu(\mathbf{P})\,d_{\min}}\sqrt{\frac{\max\{m, \log(n)\}}{L}}\,,$$

where $\kappa = \log\left(\frac{d_{\text{avg}}}{d_{\min} w_{\min}}\right)$, $w_{\min} = \min_{i\in[n]} w_i$, $d_{\text{avg}} = \sum_{i\in[n]} w_i d_i$, $d_{\min} = \min_{i\in[n]} d_i$, $\mu(\mathbf{P})$ is the spectral gap of the random walk $\mathbf{P}$ (Eq. (1)), and C is any constant.

Let us first state the concentration inequality for multinomial distributions due to (Devroye, 1983), which will be useful in proving this theorem.

Lemma 4 (Multinomial distribution inequality)**.** *(Devroye, 1983) Let $Y_1, \ldots, Y_n$ be a sequence of n independent random variables drawn from the multinomial distribution with parameters $(p_1, \ldots, p_k)$. Let X_i be the number of times i occurs in the n draws, i.e. $X_i = \sum_{j=1}^n \mathbf{1}[Y_j = i]$. For all $\epsilon \in (0,1)$, and all k satisfying $k/n \leq \epsilon^2/20$, we have*

$$P(\sum_{i=1}^k |X_i - np_i| \geq n\epsilon) \leq 3\exp(-n\epsilon^2/25).$$

To prove Theorem 1, we shall first prove a bound on the total variation distance between the stationary states $\boldsymbol{\pi}$ and $\hat{\boldsymbol{\pi}}$ of the transition matrices $\mathbf{P}$ and $\hat{\mathbf{P}}$ respectively. We shall then prove a bound on the distance between the true weights $\mathbf{w}$ and estimates $\hat{\mathbf{w}}$ in terms of the distance between $\boldsymbol{\pi}$ and $\hat{\boldsymbol{\pi}}$.

An important result in the stability theory of Markov chains shows a connection between the stability of a chain and its speed of convergence to equilibrium (Mitrophanov, 2005). In fact, we can bound the sensitivity of a Markov chain under perturbation as a function of the convergence rate of the chain, with the accuracy of the sensitivity bound depending on the sharpness of the bound on the convergence rate. The following theorem is a specialization of Theorem 3.1 of (Mitrophanov, 2005), which gives perturbation bounds for Markov chains with general state spaces.

Theorem 2. *(Mitrophanov, 2005) Consider two discrete-time Markov chains $\mathbf{P}$ and $\hat{\mathbf{P}}$, with finite state space $\Omega = \{1, \ldots, n\}, n \geq 1$, and stationary distributions $\boldsymbol{\pi}$ and $\hat{\boldsymbol{\pi}}$, respectively. If there exist positive constants $1 < R < \infty$ and $\rho < 1$ such that*

$$\max_{x\in\Omega} \|\mathbf{P}^t(x,\cdot) - \boldsymbol{\pi}\|_{TV} \leq R\rho^t, \qquad \forall t \in \mathbb{N}$$

then for $\mathbf{E} := \mathbf{P} - \hat{\mathbf{P}}$, we have

$$\|\boldsymbol{\pi} - \hat{\boldsymbol{\pi}}\|_{TV} \leq \left(\hat{t} + \frac{1}{1-\rho}\right) \cdot \|\mathbf{E}\|_\infty\,.$$

where $\hat{t} = \log(R)/\log(1/\rho)$, and $\|\cdot\|_\infty$ is the matrix norm induced by the L_∞ vector norm.

It is well known that all ergodic Markov chains satisfy the conditions imposed by Theorem 2. In order to obtain sharp bounds on the convergence rate, we shall leverage the fact that the (unperturbed) Markov chain corresponding to the ideal transition probability matrix $\mathbf{P}$ is time-reversible.

Theorem 3. *(Diaconis & Stroock, 1991) Let $\mathbf{P}$ be an irreducible, reversible Markov chain with finite state space $\Omega = \{1, \ldots, n\}, n \geq 1$, and stationary distribution $\boldsymbol{\pi}$. Let $\lambda_2 := \lambda_2(\mathbf{P})$ be the second largest eigenvalue of $\mathbf{P}$ in terms of absolute value. Then for all $x \in \Omega$, $t \in \mathbb{N}$,*

$$\|\mathbf{P}^t(x,\cdot) - \boldsymbol{\pi}\|_{TV} \leq \sqrt{\frac{1-\pi(x)}{4\pi(x)}}\lambda_2^t$$

Comparing these bounds with the conditions imposed by

Theorem 2, we can observe that

$$\begin{aligned}\rho &= \lambda_2, \\ R &= \max_{i\in[n]}\sqrt{\frac{1-\pi(i)}{4\pi(i)}} \\ &= \max_{i\in[n]}\sqrt{\frac{\|\mathbf{D}\mathbf{w}\|_1 - w_i d_i}{4 w_i d_i}} \\ &\le \sqrt{\frac{d_{\text{avg}}}{4 d_{\min} w_{\min}}}\,,\end{aligned}$$

where $w_{\min} = \min_{i\in[n]} w_i$. Substituting these values into the perturbation bounds of Theorem 2, we get

$$\begin{aligned}\hat{t} + \frac{1}{1-\rho} &= \frac{\log(d_{\text{avg}}/(4d_{\min}w_{\min}))}{2\log(1/\lambda_2(\mathbf{P}))} + \frac{1}{1-\lambda_2(\mathbf{P})} \\ &\le \frac{\log(d_{\text{avg}}/(4d_{\min}w_{\min}))}{2(1-\lambda_2(\mathbf{P}))} + \frac{1}{1-\lambda_2(\mathbf{P})} \\ &< \frac{\kappa}{2\mu(\mathbf{P})}, \qquad \text{where } \kappa = \log(\frac{2d_{\text{avg}}}{d_{\min}w_{\min}})\end{aligned}$$

Now, the next step is to show that the perturbation error $\mathbf{E} := \mathbf{P} - \hat{\mathbf{P}}$ is bounded in terms of the matrix L_∞ norm.

Lemma 5. *For $\mathbf{E} := \mathbf{P} - \hat{\mathbf{P}}$, we have with probability $\ge 1 - 3n^{-(C^2-50)/25}$,*

$$\|\mathbf{E}\|_\infty \le C\sqrt{\frac{\max\{m, \log n\}}{L}}$$

where C is any constant.

Proof. By definition, $\|\mathbf{E}\|_\infty = \max_i \sum_{j=1}^n |\hat{P}_{ij} - P_{ij}|$. Fix any row $i \in [n]$. The probability that the absolute row sum exceeds a fixed positive quantity t is given by

$$\begin{aligned}&P(\sum_{j=1}^n |\hat{P}_{ij} - P_{ij}| \ge t) \\ &= P(\sum_{j=1}^n |\frac{1}{d_i}\sum_{a:i,j\in S_a}(\hat{p}_{j|S_a} - p_{j|S_a})| \ge t) \\ &= P(\sum_{j=1}^n |\frac{1}{d_i}\sum_{a:i,j\in S_a}\frac{1}{L}\sum_{l=1}^L(\mathbf{1}(y_a^l = j) - p_{j|S_a})| \ge t) \\ &\le P(\sum_{j=1}^n \sum_{a:i,j\in S_a} |\sum_{l=1}^L(\mathbf{1}(y_a^l = j) - p_{j|S_a})| \ge L d_i t) \\ &= P(\sum_{a:i\in S_a}\sum_{j\in S_a} |\sum_{l=1}^L(\mathbf{1}(y_a^l = j) - p_{j|S_a})| \ge L d_i t) \\ &\le d_i P(\sum_{j\in S_a} |\sum_{l=1}^L(\mathbf{1}(y_l^a = j) - p_{j|S_a})| \ge \frac{L d_i t}{d_i})\end{aligned}$$

with the final pair of inequalities following from rearranging the terms in the summations and applying union bound. We leverage the multinomial distribution concentration inequality (Lemma 4) of Devroye (1983) to obtain the following bound for any set S_a for any m satisfying a technical condition $m/L \le t^2/20$.

$$P(\sum_{j\in S_a} |\sum_{l=1}^L(\mathbf{1}(y_l^a = j) - p_{j|S_a})| \ge Lt) \le 3\exp(\frac{-Lt^2}{25})$$

Thus, using union bound, the probability that any absolute row sum exceeds t is at most $3nd_{\max}\exp(-Lt^2/25)$. By selection of $t = 5C'\sqrt{\max\{m, \log n\}/L}$, we get

$$\begin{aligned}&P\left(\|\mathbf{E}\|_\infty \ge 5C'\sqrt{\frac{\max\{m,\log n\}}{L}}\right) \\ &\le 3n^2\exp\left(\frac{-25C'^2 L\max\{m,\log n\}}{25L}\right) \\ &\le 3n^{-(C'^2-2)}\end{aligned}$$

substituting $C = 5C'$ proves our claim. Lastly, one can verify that the aforementioned choice of t satisfies the technical condition imposed by Lemma 4 for any n, m and L. □

Combining the results of Theorem 2, Theorem 3, and Theorem 5 gives us a high confidence total variation error bound on the stationary states $\boldsymbol{\pi}$ and $\hat{\boldsymbol{\pi}}$ of the ideal and perturbed Markov chains $\mathbf{P}$ and $\hat{\mathbf{P}}$ respectively. Thus, with confidence $\ge 1 - 3n^{-(C^2-50)/25}$, we have

$$\|\boldsymbol{\pi} - \hat{\boldsymbol{\pi}}\|_{\text{TV}} \le \frac{C\kappa}{\mu(\mathbf{P})}\sqrt{\frac{\max\{m,\log n\}}{L}}\,, \tag{9}$$

where $\kappa = \log(2d_{\text{avg}}/(d_{\min}w_{\min}))$.

The last step in our scheme is to prove that the linear transformation $\mathbf{D}^{-1}\hat{\boldsymbol{\pi}}$ preserves this error bound up to a reasonable factor.

Lemma 6. *Under the conditions of Theorem 1, let $\boldsymbol{\pi} = \mathbf{D}\mathbf{w}/\|\mathbf{D}\mathbf{w}\|_1$ and $\hat{\boldsymbol{\pi}} = \mathbf{D}\hat{\mathbf{w}}/\|\mathbf{D}\hat{\mathbf{w}}\|_1$ be the unique stationary distributions of the Markov chains $\mathbf{P}$ (Eq. (1)) and $\hat{\mathbf{P}}$ (Eq. (2)) respectively. Then we have*

$$\|\mathbf{w} - \hat{\mathbf{w}}\|_{TV} \le \frac{d_{\text{avg}}}{d_{\min}}\|\boldsymbol{\pi} - \hat{\boldsymbol{\pi}}\|_{TV}\,.$$

Proof. We shall divide our proof into two cases.
Case 1: $\|\mathbf{D}\hat{\mathbf{w}}\|_1 \ge \|\mathbf{D}\mathbf{w}\|_1$.
Let us define the set $A = \{i : w_i \ge \hat{w}_i\}$, and the set $A' = \{j : \pi_j \ge \hat{\pi}_j\}$. When $\|\mathbf{D}\hat{\mathbf{w}}\|_1 \ge \|\mathbf{D}\mathbf{w}\|_1$, it is easy to see that $A \subseteq A'$.

Consider the total variation distance $\|\mathbf{w}-\hat{\mathbf{w}}\|_{TV}$ between the true preferences $\mathbf{w}$ and our estimates $\hat{\mathbf{w}}$. By definition,

$$\begin{aligned}
\|\mathbf{w}-\hat{\mathbf{w}}\|_{TV} &= \sum_{i\in A}(w_i-\hat{w}_i) \\
&= \sum_{i\in A} w_i\left(1-\frac{\hat{w}_i}{w_i}\right) = \sum_{i\in A} w_i\left(1-\frac{\hat{w}_i d_i}{w_i d_i}\right) \\
&\le \sum_{i\in A} w_i\left(1-\frac{\hat{w}_i d_i\|\mathbf{D}\mathbf{w}\|_1}{w_i d_i\|\mathbf{D}\hat{\mathbf{w}}\|_1}\right) \\
&= \sum_{i\in A} w_i\left(1-\frac{\hat{\pi}_i}{\pi_i}\right) \\
&= \sum_{i\in A} w_i\left(\frac{(\pi_i-\hat{\pi}_i)\|\mathbf{D}\mathbf{w}\|_1}{w_i d_i}\right) \\
&\le \sum_{j\in A'} w_j\left(\frac{(\pi_j-\hat{\pi}_j)\|\mathbf{D}\mathbf{w}\|_1}{w_j d_j}\right) \\
&= \sum_{j\in A'}\left(\frac{(\pi_j-\hat{\pi}_j)\|\mathbf{D}\mathbf{w}\|_1}{d_j}\right) \\
&\le \frac{\|\mathbf{D}\mathbf{w}\|_1}{d_{\min}}\sum_{j\in A'}(\pi_j-\hat{\pi}_j) = \frac{d_{\text{avg}}}{d_{\min}}\|\pi-\hat{\pi}\|_{TV}
\end{aligned}$$

Case 2, where $\|\mathbf{D}\hat{\mathbf{w}}\|_1 < \|\mathbf{D}\mathbf{w}\|_1$ follows symmetrically, giving us the inequality

$$\begin{aligned}
\|\mathbf{w}-\hat{\mathbf{w}}\|_{TV} &\le \frac{\|\mathbf{D}\hat{\mathbf{w}}\|_1}{d_{\min}}\|\pi-\hat{\pi}\|_{TV} \\
&\le \frac{\|\mathbf{D}\mathbf{w}\|_1}{d_{\min}}\|\pi-\hat{\pi}\|_{TV} = \frac{d_{\text{avg}}}{d_{\min}}\|\pi-\hat{\pi}\|_{TV}
\end{aligned}$$

where the last inequality follows from the assumption of Case 2, proving our claim. □

Combining the above lemma with Eq. (9) gives us the statement of the theorem.

E. Proof of Corollary 1

Corollary 1. *In the setting of Theorem 1, the* ASR *algorithm converges to a unique distribution $\hat{\mathbf{w}}$, which with probability $\ge 1-3n^{-(C^2-50)/25}$ satisfies the following error bound:*

$$\|\mathbf{w}-\hat{\mathbf{w}}\|_{TV} \le \frac{C\, m\, b^2\, \kappa\, d_{\text{avg}}}{\xi\, d_{\min}}\sqrt{\frac{\max\{m,\log(n)\}}{L}}\,,$$

where $b = \max_{i,j\in[n]}\frac{w_i}{w_j}$.

Corollary 1 follows from the following lemma which compares the spectral gap of the matrix $\mathbf{P}$ with the spectral gap of the graph Laplacian.

Lemma 7. *Let $\mathbf{L} := \mathbf{C}^{-1}\mathbf{A}$ be the Laplacian of the undirected graph $G_c([n],E)$. Then the spectral gap $\mu(\mathbf{P}) = 1-\lambda_2(\mathbf{P})$ of the reversible Markov chain $\mathbf{P}$ (Eq. (2)) corresponding to the* ASR *algorithm is related to the spectral gap $\xi = 1-\lambda_2(\mathbf{L})$ of the Laplacian as*

$$\mu(\mathbf{P}) \ge \frac{\xi}{mb^2}$$

Proof. To prove this inequality, we shall leverage the comparison Lemma 3 of (Diaconis & Saloff-Coste, 1993), with $\mathbf{Q},\boldsymbol{\nu} = \mathbf{L},\boldsymbol{\nu}$. From the definition of the Laplacian, it is clear that for all i, $\nu_i\mathbf{L}_{ij} = 1/2|E|$. Furthermore, $\nu_i = c_i/2|E| \ge d_i/2|E|$, where c_i is the number of unique items i was compared with, which is trivially at least the number of unique multiway comparisons of which i was a part. Thus,

$$\begin{aligned}
\beta &:= \max_{i\in[n]}\frac{\pi_i}{\nu_i} = \max_{i\in[n]}\frac{w_i d_i/\|\mathbf{D}\mathbf{w}\|_1}{c_i/2|E|} \\
&\le \frac{2|E|w_{\max}}{\|\mathbf{D}\mathbf{w}\|_1} \\
\alpha &:= \min_{(i,j)\in E}\frac{\pi_i P_{ij}}{\nu_i L_{ij}} \\
&= \min_{(i,j)\in E}\frac{\frac{w_i d_i}{\|\mathbf{D}\mathbf{w}\|_1}\frac{1}{d_i}\sum_{a:(i,j)\in S_a}\frac{w_j}{\sum_{k\in S_a}w_k}}{1/2|E|} \\
&\ge \frac{2|E|w_{\min}^2}{m w_{\max}\|\mathbf{D}\mathbf{w}\|_1}
\end{aligned}$$

Thus, $\alpha/\beta \ge 1/mb^2$, which proves our claim. □

F. Proof of Corollary 2

Corollary 2. *If the conditions of Theorem 1 are satisfied, and if the number of edges in the comparison graph G_c are $O(n\,\text{poly}(\log n))$, i.e. $|E| = O(n\,\text{poly}(\log n))$, then in order to ensure a total variation error of $o(1)$, the required number of comparisons per set is upper bounded as*

$$L = O\big(\mu(\mathbf{P})^{-2}\,\text{poly}(\log n)\big) = O\big(\xi^{-2}\, m^3\,\text{poly}(\log n)\big)\,.$$

Hence, the sample complexity, i.e. total number of m-way comparisons needed to estimate $\mathbf{w}$ with error $o(1)$, is given by $|E|\times L = O\big(\xi^{-2}\, m^3\, n\,\text{poly}(\log n)\big)$.

In order to prove the above corollary we first give the following claim.

Claim 1. *Given items $[n]$, and comparison graph $G_c = ([n],E)$ induced by comparison data $\mathbf{Y} = \{S_a,\mathbf{y}_a\}_{a=1}^d$, let the vector of true MNL parameters be $\mathbf{w} = (w_1,\ldots,w_n)$. Furthermore, let d_i represent the number of unique comparisons of which item $i\in[n]$ was a part. Then we have*

$$d_{\text{avg}} = \sum_{i\in[n]} w_i d_i \le \frac{2w_{\max}|E|}{w_{\min}n}\,,$$

where $w_{\max} = \max_{i\in[n]} w_i$, *and* $w_{\min} = \min_{j\in[n]} w_j$.

Proof. Clearly,

$$w_{\min} \sum_{i\in[n]} w_i d_i \le \frac{1}{n} \sum_{i\in[n]} w_i d_i \le \frac{w_{\max}}{n} \sum_{i\in[n]} d_i,$$

The statement of the lemma follows by realizing that $\sum_{i\in[n]} d_i \le \sum_{i\in[n]} c_i \le 2|E|$. □

Proof. (of Corollary 2) Substituting the above bound on d_{avg} in the sample complexity bounds of Corollary 1, we get the following guarantee on the total variation error between the estimates $\hat{\mathbf{w}}$ and the true weight vector $\mathbf{w}$

$$\|\mathbf{w} - \hat{\mathbf{w}}\|_{\text{TV}} \le \frac{C\, m\, b^3\, \kappa\, |E|}{n\, \xi\, d_{\min}} \sqrt{\frac{\max\{m, \log(n)\}}{L}}\,,$$

where $b = \frac{w_{\max}}{w_{\min}}$. Furthermore, this guarantee holds with probability $\ge 1 - 3n^{-(C^2-50)/25}$. From this, we can conclude that if

$$L \ge \max\{m, \log(n)\} \left(\frac{10\, m\, b^3\, \kappa\, |E|}{n\, \xi\, d_{\min}}\right)^2,$$

then it is sufficient to guarantee that $\|\mathbf{w} - \hat{\mathbf{w}}\|_{\text{TV}} = o(1)$ with probability $\ge 1 - 3n^{-2}$. Trivially bounding $\kappa = O(\log n)$, and from the assumptions $b = O(1)$ and $|E| = O(n\,\text{poly}(\log n))$, we can conclude

$$L = O(\xi^{-2} m^3 \text{poly}(\log n))$$

where the additional m factor comes from trivially bounding $\max\{m, \log n\} \le m \log n$. This gives us a sample complexity bound of

$$|E| \times L = O(\xi^{-2}\, m^3\, n\, \text{poly}(\log n))$$

for our algorithm, which proves the corollary. □

G. Proof of Lemma 2

Lemma 2. *For any realization of comparison data* $\mathbf{Y}$*, there is a one-to-one correspondence d each iteration of the message passing algorithm (2) and the corresponding power iteration of the ASR algorithm (1), and both algorithms return the same estimates* $\hat{\mathbf{w}}$ *for any* $\mathbf{Y}$.

Proof. In the message passing algorithm, the item to set messages $m^{(r)}_{i\to a}$ in round r correspond to the estimates of the item weights. One can verify that the estimate $\hat{w}^{(r)}$ of item i in round r evolves according to the following equation.

$$\hat{w}_i^{(r+1)} = \frac{1}{d_i} \sum_{a:i\in S_a} p_{i|S_a} \cdot \sum_{j\in S_a} \hat{w}_j^{(r)}\,.$$

We can represent this system of equations compactly using the following matrices. Let $\hat{\mathbf{V}} \in \mathbb{R}^{d\times n}$ be a matrix such that

$$\hat{V}_{ai} := \begin{cases} \frac{p_{i|S_a}}{d_i} & \text{if } (i,a) \in E \\ 0 & \text{otherwise} \end{cases}, \tag{10}$$

and $\mathbf{B} \in \mathbb{R}^{n\times d}$ be a matrix such that

$$B_{ia} := \begin{cases} 1 & \text{if } (i,a) \in E \\ 0 & \text{otherwise} \end{cases}, \tag{11}$$

Thus, we can represent the weight update from round (r) to round $(r+1)$ as

$$\begin{aligned} \hat{\mathbf{w}}^{(r+1)} &= (\mathbf{B}\hat{\mathbf{V}})^\top \hat{\mathbf{w}}^{(r)} = \hat{\mathbf{M}}^\top \hat{\mathbf{w}}^{(r)} \\ &= (\hat{\mathbf{M}}^\top)^r \hat{\mathbf{w}}^{(0)}\,, \end{aligned}$$

where $\hat{\mathbf{M}} := \mathbf{B}\hat{\mathbf{V}}$, with entry (i,j) of $\hat{\mathbf{M}}$ being

$$\hat{M}_{ij} := \frac{1}{d_j} \sum_{a:i,j\in S_a} p_{j|S_a}\,. \tag{12}$$

The above equation implies that the message passing algorithm is essentially a power iteration on the matrix $\hat{\mathbf{M}}$. Now, it is easy to see that $\hat{\mathbf{M}} = \mathbf{D}\hat{\mathbf{P}}\mathbf{D}^{-1}$ where $\hat{\mathbf{P}}$ is the transition matrix constructed by ASR ($Eq.$ (2)). Therefore, there is a one-to-one correspondence between the power iterations on $\hat{\mathbf{M}}$ and $\hat{\mathbf{P}}$. More formally, if we initialize with $\hat{\mathbf{w}}^{(0)}$ in the power iteration on $\hat{\mathbf{M}}$, and initialize with $\hat{\boldsymbol{\pi}}^{(0)} = \mathbf{D}\hat{\mathbf{w}}^{(0)}$ in the power iteration on $\mathbf{P}$, then the iterates at the r-th step will be related as $\hat{\boldsymbol{\pi}}^{(r)} = \mathbf{D}\hat{\mathbf{w}}^{(r)}$. Furthermore, if $\hat{\boldsymbol{\pi}}$ is the stationary distribution of $\hat{\mathbf{P}}$, then $\hat{\mathbf{w}} = \mathbf{D}^{-1}\hat{\boldsymbol{\pi}}$ is the corresponding dominant left eigenvector of $\hat{\mathbf{M}}$, i.e. $\mathbf{D}^{-1}\hat{\boldsymbol{\pi}} = \hat{\mathbf{M}}^\top\mathbf{D}^{-1}\hat{\boldsymbol{\pi}}$. Also, $\hat{\mathbf{w}}$ is exactly the estimate (after normalization) returned by both the ASR and the message passing algorithm upon convergence. Thus, we can conclude that the message passing algorithm is identical to ASR for any realization of comparison data generated according to the MNL model. □

H. Additional Experimental Results

In this section we will describe additional experimental results comparing our algorithm and the RC/LSR algorithms on various synthetic and real world datasets. Since we require additional regularization when the random walk induced by comparison data is reducible, we will first describe the regularized version of the RC and LSR algorithms (regularized version of our algorithm is given in Appendix A).

H.1. RC and LSR algorithms with regularization

In this section, for the sake of completeness, we state the regularized version of the RC (Negahban et al., 2017) and

Table 2. Statistics for real world datasets

Dataset	n	m	d	**total choices**
Youtube	21207	2	394007	1138562
GIF-amusement	6118	2	75649	77609
GIF-anger	6119	2	64830	66505
GIF-contentment	6118	2	70230	72175
GIF-excitement	6119	2	80493	82564
GIF-happiness	6119	2	104801	107816
GIF-pleasure	6119	2	86499	88959
GIF-relief	6112	2	38770	39853
GIF-sadness	6118	2	63577	65263
GIF-satisfaction	6118	2	78401	80474
GIF-shame	6116	2	46249	47550
GIF-surprise	6118	2	63850	65591
SFWork	6	3-6	12	5029
SFShop	8	4-8	10	3157

LSR (Maystre & Grossglauser, 2015) algorithms.[7] These algorithms are based on computing the stationary distribution of a Markov chain. In the case of pairwise comparisons, for a regularization parameter $\lambda > 0$, the Markov chain $\hat{\mathbf{P}}'^{\text{RC}} := [\hat{P}'^{\text{RC}}_{ij}]$, where, $\forall i, j \in [n]$,

$$\hat{P}'^{\text{RC}}_{ij} := \begin{cases} \frac{1}{d_{\max}} \left(\frac{n_{j|\{i,j\}}+\lambda}{n_{j|\{i,j\}}+n_{i|\{i,j\}}+2\lambda} \right), & \text{if } i \neq j \\ 1 - \frac{1}{d_{\max}} \sum_{j' \neq i} \hat{P}'^{\text{RC}}_{ij}, & \text{if } i = j \end{cases}$$

and $n_{j|\{i,j\}}$ is defined according to Eq. (7). In the case of multi-way comparisons, the Markov chain $\hat{\mathbf{P}}'^{\text{LSR}} := [\hat{P}'^{\text{LSR}}_{ij}]$, where, $\forall i, j \in [n]$,

$$\hat{P}'^{\text{LSR}}_{ij} := \begin{cases} \epsilon \sum_{a \in [d]: i,j \in S_a} \left(\frac{n_{j|S_a}+\lambda}{|S_a|} \right), & \text{if } i \neq j \\ 1 - \epsilon \sum_{j' \neq i} \hat{P}'^{\text{LSR}}_{ij}, & \text{if } i = j \end{cases}$$

where ϵ is a quantity small enough to make the diagonal entries of $\hat{\mathbf{P}}'^{\text{LSR}}$ non negative, and $n_{j|S_a}$ is again defined according to Eq. (7).

H.2. Synthetic Datasets

In this section, we give additional experimental results for various other values of parameters m and n. The plots are given in the figures below. The general trends observed from these experiments are exactly as predicted by our theoretical analysis. In particular, we note that even in the case of a star graph topology, the convergence rate of ASR remains essentially the same with increasing n, while the performance of RC and LSR degrades smoothly. This really conveys the low dependence on the ratio $d_{\max}/d_{\min}$.

[7] See Section 3.3 in Negahban et al. (2017) for more details.

H.3. Real Datasets

In this section, we provide additional experimental results for more datasets, and additional values of the regularization parameter λ. We conducted experiments on the YouTube dataset (Shetty, 2012), various GIF datasets (Rich et al.), and the SFwork and SFshop (Koppelman & Bhat, 2006) datasets. Below we briefly describe each of these datasets (additional statistics are given in Table 2).

1. **YouTube Comedy Slam Preference Data**. This dataset is due to a video discovery experiment on YouTube in which users were shown a pair of videos and were asked to vote for the video they found funnier out of the two.[8]
2. **GIFGIF datasets**. These datasets are due to a experiment that tries to understand the emotional content present in animated GIFs. In this experiment users are shown a pair of GIFs and asked to vote for the GIF that most accurately represents a particular emotion. These votes are collected for several different emotions.[9]
3. **SF datasets**. These datasets are from a survey of transportation preferences around the San Francisco Bay Area in which citizens were asked to vote on their preferred commute option amongst different options.[10]

As expected, the peak log likelihood decreases with increasing λ, as this regularization parameter essentially dampens the information imparted by the comparison data. We also plot degree distributions of these real world datasets in order to explore the behavior of the ratio $d_{\max}/d_{\min}$ in practice. In particular, we observe that this quantity does not really behave like a constant, and is very large in most cases. This is particularly evident in the Youtube dataset, where the degree distribution closely follows the power law relationship with n.

[8] See `https://archive.ics.uci.edu/ml/datasets/YouTube+Comedy+Slam+Preference+Data` for more details.

[9] See `http://gif.gf` for more details.

[10] These datasets are available at `https://github.com/sragain/pcmc-nips`.

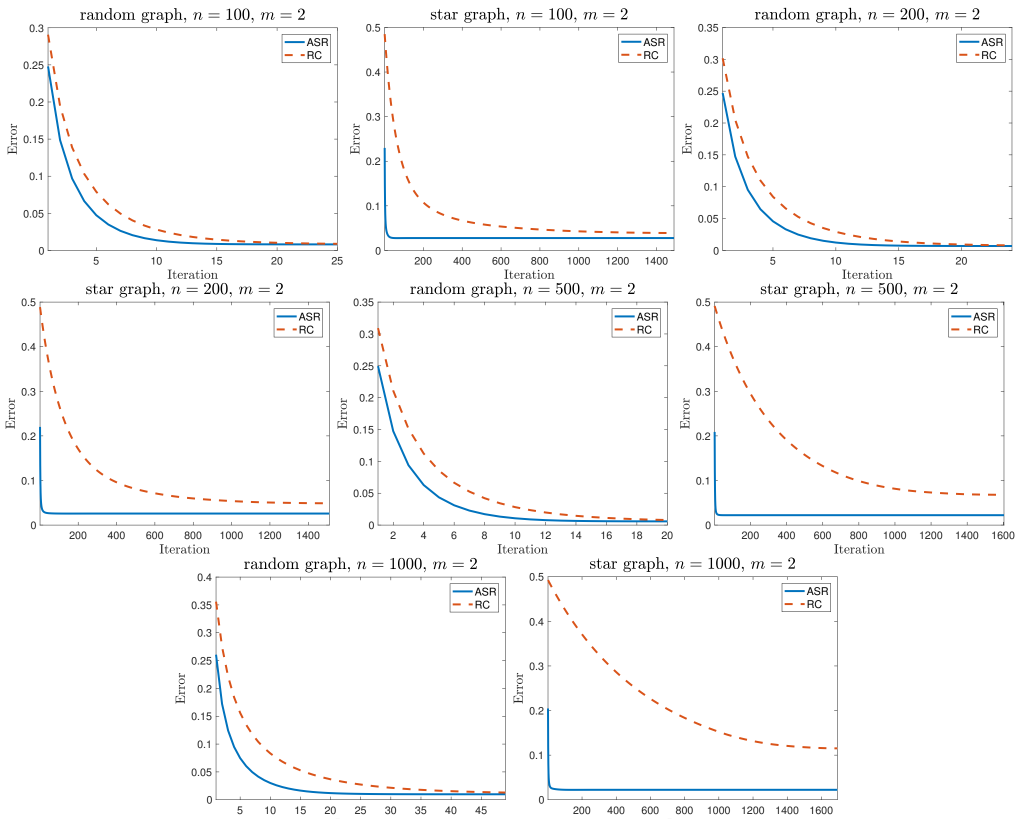

Figure 3. Results on synthetic data: L_1 error vs. number of iterations for our algorithm, ASR, compared with the RC algorithm (for $m = 2$) on data generated from the MNL/BTL model with the random and star graph topologies.

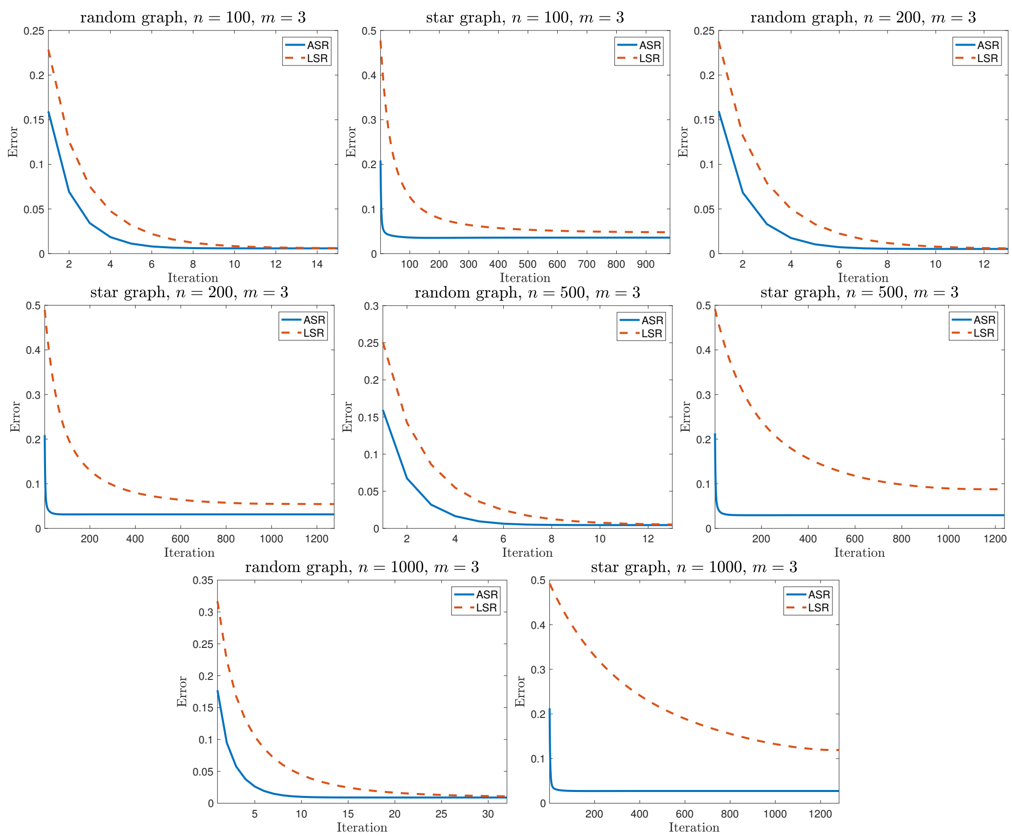

Figure 4. Results on synthetic data: L_1 error vs. number of iterations for our algorithm, ASR, compared with the LSR algorithm (for $m = 3$) on data generated from the MNL/BTL model with the random and star graph topologies.

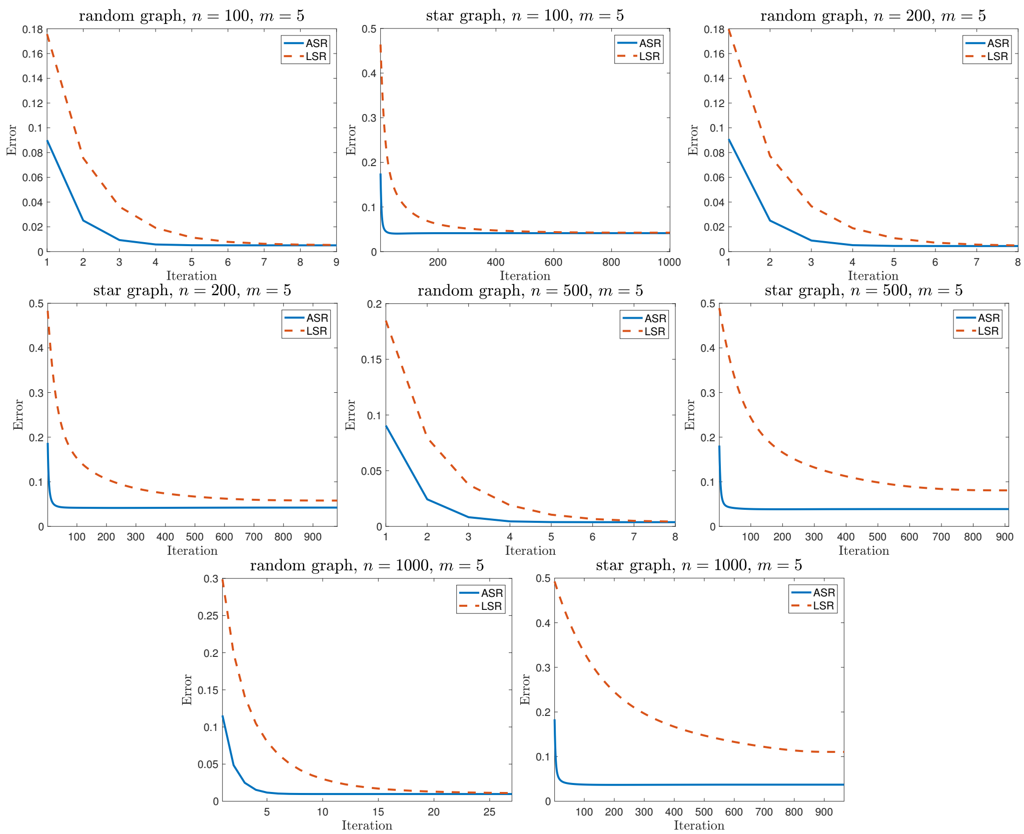

Figure 5. Results on synthetic data: L_1 error vs. number of iterations for our algorithm, ASR, compared with the LSR algorithm (for $m = 5$) on data generated from the MNL/BTL model with the random and star graph topologies.

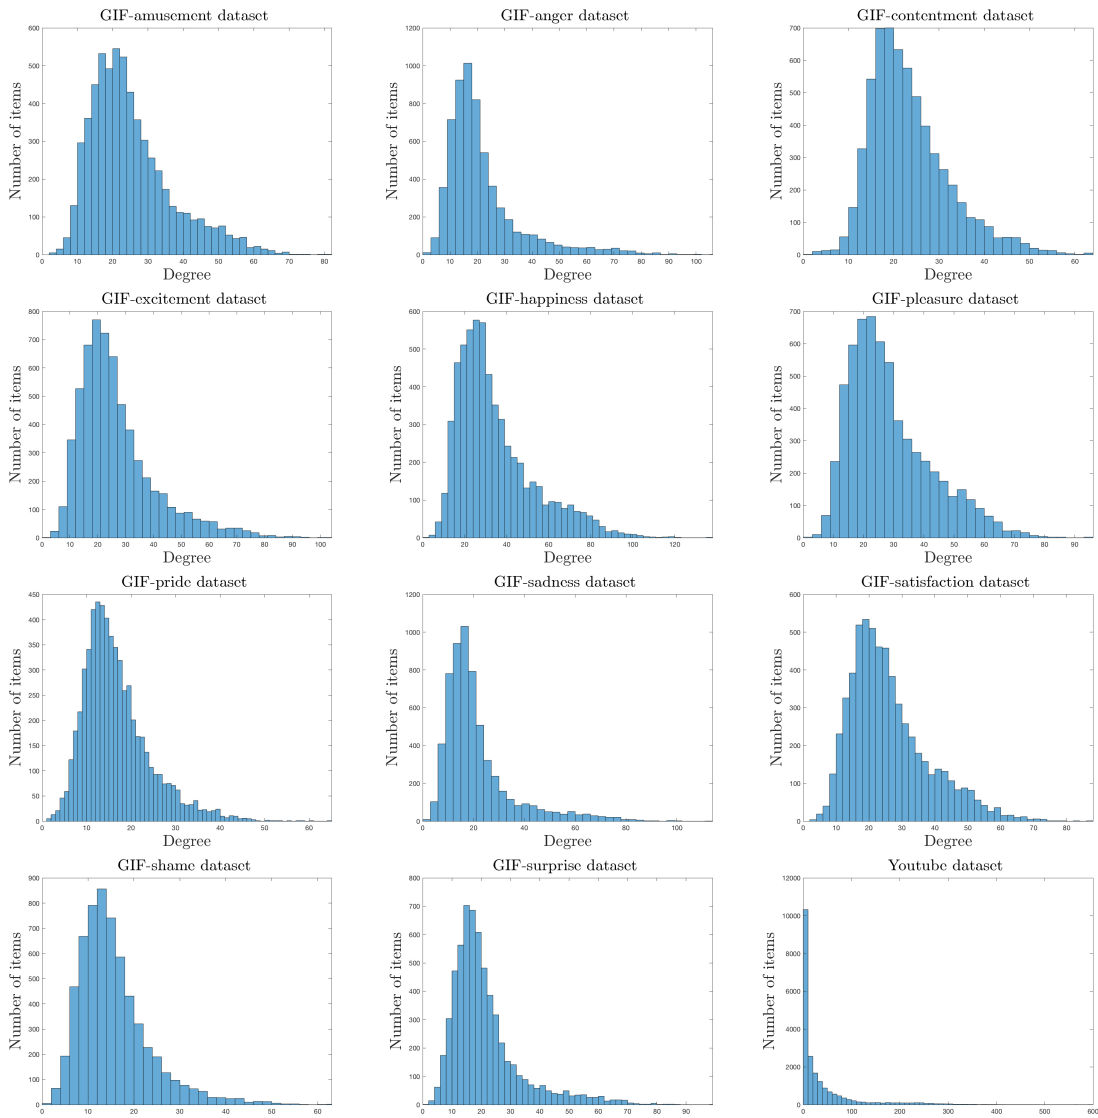

Figure 6. Degree distributions of various real world datasets.

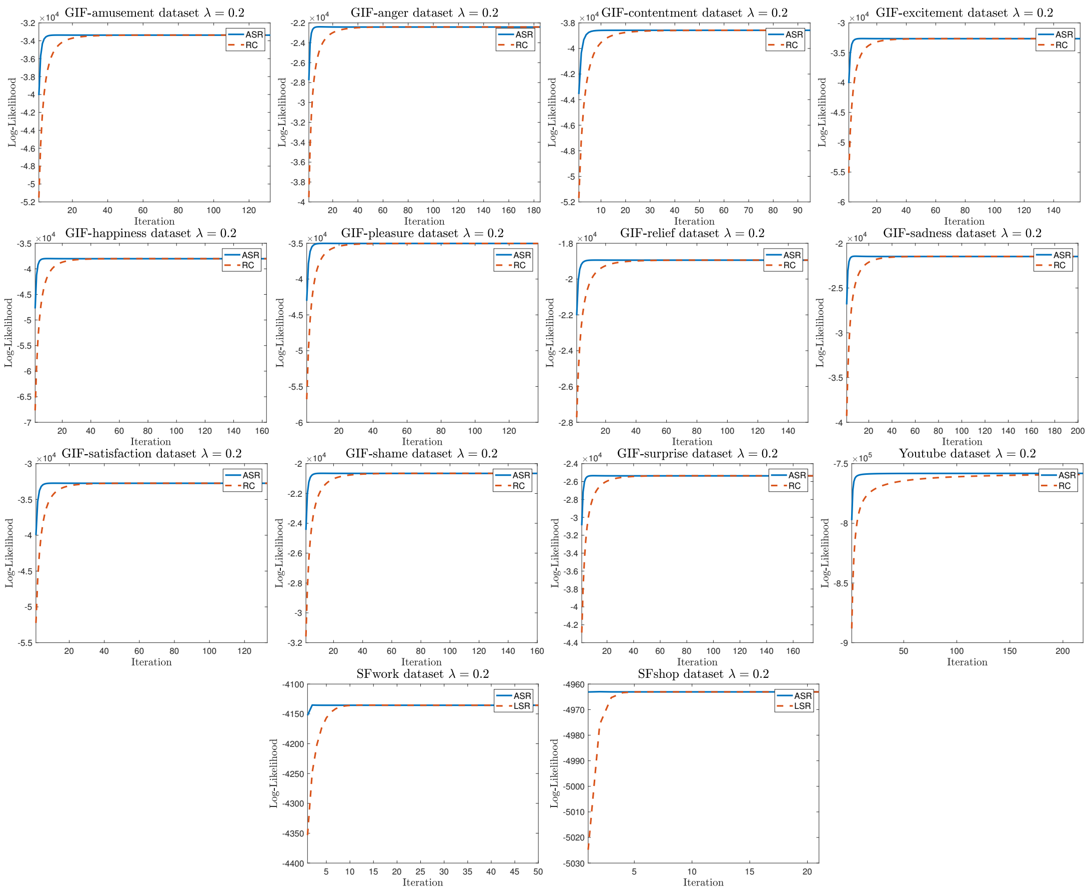

Figure 7. Results on real data: Log-likelihood vs. number of iterations for our algorithm, ASR, compared with the RC algorithm (for pairwise comparison data) and the LSR algorithm (for multi-way comparison data), all with regularization parameter set to 0.2.

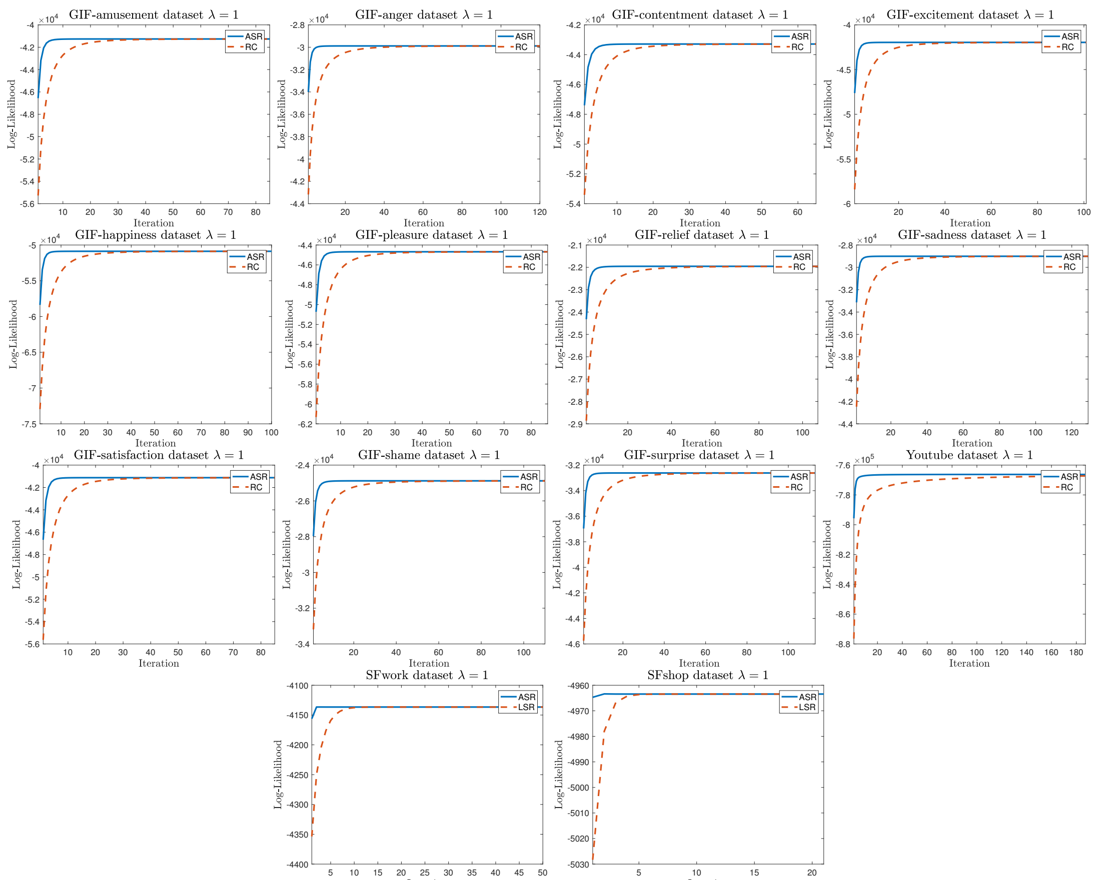

Figure 8. Results on real data: Log-likelihood vs. number of iterations for our algorithm, ASR, compared with the RC algorithm (for pairwise comparison data) and the LSR algorithm (for multi-way comparison data), all with regularization parameter set to 1.