Learning Mallows Models with Pairwise Preferences: Supplementary Material

1 Proofs

Theorem 1 ([1]). *Given a partial order v, computing the number of linear extensions of v, that is $|\Omega(v)|$, is #P-complete.*

To show that computing a function $f(x)$ is #P-hard for input x, it is sufficient to show that a #P-complete problem can be reduced to it in polynomial time.

Proof of Proposition 5. We reduce counting the number of linear extensions to this problem. Given v, notice that any $r = r_1 \dots r_m \in \Omega(v)$ has a uniform posterior probability of $1/|\Omega(v)|$. Let $\Phi_\sigma^{-1}(r) = (j_1, \dots, j_m)$. Now make $m-1$ calls to the subroutine for computing GRIM insertion probabilities p_{ij_i} with partial order v for each $i \in \{2, \dots, m\}$. The posterior probability of r is $1/|\Omega(v)| = \prod_i p_{ij_i}$ this implies we can compute $|\Omega(v)|$ by inverting the product of the insertion probabilities. Note that this reduction is polytime because we use a topological sort to find u. □

Proof of Theorem 11. We reduce counting the number of linear extensions to this problem. Let v be a partial order (i.e. an input to counting linear extensions), encode the input to the log likelihood computation as follows: let $V = (v)$, $K = 1$ with $\phi = 1$ (σ can be any ranking). Hence $\mathcal{L} = \mathcal{L}(\pi, \sigma, \phi | V) = \ln \sum_{r \in \Omega(v)} 1/m!$. Thus we can recover the number of linear extensions by computing $\exp(\mathcal{L}) \cdot m!$. We can do this in polytime by noting that $\mathcal{L}$ is polynomial in m and by using the power series expansion $\sum_{i \geq 0} \mathcal{L}^i m!/i!$ where we can simplify truncating the series at polynomial number of steps when the terms no longer impact the number's integer portion. □

Proof of Proposition 6. Inserting σ_i in any rank less than l_i is impossible since either $l_i = 1$ (can't insert in rank 0) or σ_i is above r_{l_i} which contradicts the requirement in $\mathsf{tc}(v)$ that r_{l_i} must be ranked higher. A similar argument can be made for inserting in rank below h_i since r_{h_i} needs to be below σ_i. Finally, inserting into any rank in $\{l_i, \dots, h_i\}$ does not violate $\mathsf{tc}(v)$ since the item will be inserted below *all* items that must precede it in $\mathsf{tc}(v)$ and *all* items that must succede it. □

Proposition 2. *For all $i \geq 2$ and all rankings of items $\sigma_1, \dots, \sigma_{i-1}$ that is consistent with v, we have that $l_i \leq h_i$. That is,* AMP *always has a position to insert item σ_i.*

Proof. Let r be a ranking of $\sigma_1, \dots, \sigma_{i-1}$ consistent with v. Let x be the lowest ranking item in r such that $x \succ_{\mathsf{tc}(v)} \sigma_i$ and y the highest ranking item in r with $y \prec_{\mathsf{tc}(v)} \sigma_i$. Thus by transitivity, $x \succ_{\mathsf{tc}(v)} y$. Now if $h_i < l_i$ (as defined in terms of r) this implies $y \succ_r x$, but this contradicts the assumption that r is consistent with v. □

Proof of Proposition 7. Since the algorithm never violates the constraints in $\mathsf{tc}(v)$, and it will always have non-empty valid insertion positions as given by Proposition 2, the algorithm will always output a ranking consistent with v. For the other direction, let $r \in \Omega(v)$ and $\Phi_\sigma^{-1}(r) = (j_1, \dots, j_m)$ the insertion ranks. We argue that for all $i \in [m]$, $j_i \in \{l_i, \dots, h_i\}$. Suppose this is not true, then there exists a smallest $i' \in [m]$ (note $i' \geq 2$ since the first item is always inserted into the first position) such that $j_{i'} \notin \{l_{i'}, \dots, h_{i'}\}$ but then by our observations about l_i and h_i this would lead to a ranking inconsistent with v—so this is not possible. Since AMP puts positive probability for any insertion position in $\{l_i, \ldots, h_i\}$ then r has positive probability under AMP. $\square$

Proof of Proposition 8. Let $\Phi_\sigma^{-1}(r) = (j_1, \ldots, j_m)$ be the insertion ranks. We have already established that AMP puts positive probability on these valid insertion ranks. In fact the probability of r under AMP is

$$\begin{aligned}\prod_{i=1}^{m} \frac{\phi^{i-j_i}}{(\phi^{i-l_i} + \phi^{i-l_i-1} + \cdots + \phi^{l-h_i})} &= \frac{\phi^{\sum_{i=1}^{m} i-j_i}}{\prod_{i=1}^{m}(\phi^{i-l_i} + \phi^{i-l_i-1} + \cdots + \phi^{l-h_i})} \\ &= \frac{\phi^{d(r,\sigma)}}{\prod_{i=1}^{m}(\phi^{i-l_i} + \phi^{i-l_i-1} + \cdots + \phi^{l-h_i})},\end{aligned}$$

where the last inequality comes from a property of the Kendall-tau metric. $\square$

Proposition 3 ([2]). *Let σ be a reference ranking. Let v be a partitioned preference with partition $A_1, \ldots, A_q$ of A. Let $\delta = \#\{(x, y) | y \succ_\sigma x, x \in A_i, y \in A_j, i, j \in [q], i < j\}$, which is the number of pairs of items across subsets of the partition that are misordered w.r.t. σ. Then*

$$\delta = \sum_{i=1}^{q-1} \sum_{x \in A_i} \sum_{j=i+1}^{q} \sum_{y \in A_j} \mathbf{1}[y \succ_\sigma x], \tag{1}$$

$$\sum_{r \in \Omega(v)} \phi^{d(r,\sigma)} = \phi^{\delta} \prod_{i=1}^{q} \prod_{j=1}^{|A_i|} (1 + \phi + \phi^2 + \cdots + \phi^{j-1}). \tag{2}$$

Proof of Proposition 9. Since the numerator of $\hat{P}_v$ part of the probability of AMP outputting r is the same as the proportional probability of the Mallows posterior, it is sufficient to show that the denominator of $\hat{P}_v$ equals the Mallows posterior normalization constant given in Eq. 2. Suppose $\sigma = \sigma_1 \cdots \sigma_m$. Consider items in A_i such that $\sigma|_{A_i} = \sigma_{t_1} \sigma_{t_2} \cdots \sigma_{t_{|A_i|}}$ (this is the ranking of items in A_i according to σ). Suppose the items $S = \{\sigma_1, \ldots, \sigma_{t_k - 1}\}$ are inserted. The structure of the resulting ranking is as follows, the items $(A_1 \cup A_2 \cup \cdots \cup A_{i-1}) \cap S$ must be in the top of the ranking, then items $A_i \cap S = \{\sigma_{t_1}, \ldots, \sigma_{t_{k-1}}\}$ are in the middle, and finally $B_{t_k} = (A_{i+1} \cup \cdots \cup A_q) \cap S$ are at bottom. When inserting σ_{t_k} into rank j, we have $j \in \{l_{t_k}, \ldots, h_{t_k}\}$ where $h_{t_k} = t_k - |B_{t_k}|$ and $l_{t_k} = h_{t_k} - |A_i \cap S| = t_k - (k-1) - |B_{t_k}|$. Hence σ_{t_k} gets inserted to rank j with probability

$$\frac{\phi^{t_k - j}}{\phi^{t_k - h_{t_k}} + \cdots + \ \phi^{t_k - l_{t_k}}} = \frac{\phi^{t_k - j}}{\phi^{|B_{t_k}|} + \cdots \phi^{k-1+|B_{t_k}|}}.$$

The denominator can be written $\phi^{|B_{t_k}|}(1 + \cdots + \phi^{k-1})$. Observe that B_{t_k} consists of all alternatives from that are above σ_{t_k} in σ, but instead are below it in v (since all these items belong to $A_{i+1} \cup \cdots \cup A_q$). So $\sum_{k=1}^{|A_i|} |B_{t_k}|$ is the total number of pairs (x, y)—where $x \in A_i$ and $y \in A_{i+1} \cup \cdots \cup A_q$—that are misordered with respect to σ. Thus inserting items in A_i will contribute a factor of

$$\prod_{k=1}^{|A_i|} \phi^{|B_{t_k}|}(1 + \cdots + \phi^{k-1}) = \phi^{\sum_{x \in A_i} \sum_{j=i+1}^{q} \sum_{y \in A_j} \mathbf{1}[y \succ_\sigma x]} \prod_{k=1}^{|A_i|} (1 + \cdots + \phi^{k-1})$$

to the denominator of $\hat{P}_v$. Once we have inserted all items, the denominator becomes

$$\phi^{\sum_{i=1}^{q} \sum_{x \in A_i} \sum_{j=i+1}^{q} \sum_{y \in A_j} \mathbf{1}[y \succ_\sigma x]} \prod_{i=1}^{q} \prod_{k=1}^{|A_i|} (1 + \cdots + \phi^{k-1}),$$

this is exactly the Mallows posterior normalization constant in Eq. 2. $\square$

Proof of Theorem 10. Note that the acceptance ratio is always positive. The proposal distribution AMP draws rankings that are independent of previous rankings and, as we proved, its support is $\Omega(v)$. Hence, for any $r' \in \Omega(v)$, MMP has positive probability of transitioning to any ranking in $\Omega(v)$ (thus establishing that $\Omega(v)$ is a recurrent class), including transitioning to itself (implying aperiodicity). □

While AMP does correspond to the Mallows posterior for the special case of partitioned preferences, in general, as we saw earlier, it won't with arbitrary paired comparisons. We will now provide some bounds on the ratio of how close the sampling algorithm is. The main technical challenge is providing a bound on the Mallows posterior normalization constant. We can get an upper bound by exploiting the paired comparison interpretation of Mallows model.

Theorem 4 (Upper Bound on Normalization)**.** *Let σ be a reference ranking, $\phi \in (0, 1]$ and v a preference. The Mallows posterior normalization constant is upper bounded by*

$$\sum_{r \in \Omega(v)} \phi^{d(r,\sigma)} \leq \phi^{d(v,\sigma)} (1 + \phi)^{\binom{m}{2} - d(v,\sigma) - s(v,\sigma)}. \tag{3}$$

where $s(v, \sigma)$ is the number of paired items in $\mathsf{tc}(v)$ that agree with σ.

Proof. We omit this lengthy proof and the required tools to a longer version of this paper. □

Eq 3 tells us if $d(v, \sigma)$ increase (i.e. v increasingly disagrees with σ) then the first factor dominates and upper bound gets smaller—this corresponds to intuition since the set $\Omega(v)$ gets "further away" from reference σ and hence its probability mass is small. Also if $|\mathsf{tc}(v)|$ is small then $d(v, \sigma) + s(v, \sigma)$ is small and the upper bound increases since the second factor will dominate. This makes sense because $\Omega(v)$ would be large and would have more probability mass. If $s(v, \sigma)$ gets larger this means more constraints in v hence $P(\Omega(v))$ would be smaller, likewise the upper bound would decrease. Before we derive a lower bound, we introduce some notions from order theory.

Definition 5. *Let v be a preference. An* anti-chain *of v is a subset X of A such that for every $x, y \in X$ they are incomparable under $\mathsf{tc}(v)$. A* maximum anti-chain *is an anti-chain whose size is at least the size of any anti-chain. The* width *of v, $\mathsf{w}(v)$ is the size of a maximum anti-chain of v.*

Theorem 6 (Lower Bound on Normalization)**.** *Let σ be a reference ranking, and $\phi \in (0, 1]$. Let X be a maximum anti-chain of v, $Y = \{a \in A \backslash X \mid \exists x \in X, a \succ_{\mathsf{tc}(v)} x\}$ and $Z = A \backslash (X \cup Y)$. Let $\delta = |\{(x, y)|x \in X, y \in Y, x \succ_\sigma y\}| + |\{(y, z)|y \in Y, z \in Z, z \succ_\sigma y\}| + |\{(x, z)|x \in X, z \in Z, z \succ_\sigma x\}|$. Denote by $\mathsf{tc}(v)|_Y$ and $\mathsf{tc}(v)|_Z$ the transitive closure of v restricted to the subsets Y and Z, respectively. Also let $\Omega(\mathsf{tc}(v)|_Y)$ denote the rankings on Y that are consistent with $\mathsf{tc}(v)|_Y$, and similarly for $\Omega(\mathsf{tc}(v)|_Z)$. We have,*

$$\sum_{r \in \Omega(v)} \phi^{d(r,\sigma)} \geq \phi^{\delta} \left[\sum_{r \in \Omega(\mathsf{tc}(v)|_Y)} \phi^{d(r,\sigma|_Y)} \right] \left[\sum_{r \in \Omega(\mathsf{tc}(v)|_Z)} \phi^{d(r,\sigma|_Z)} \right] \prod_{i=1}^{\mathsf{w}(v)} \sum_{j=0}^{i-1} \phi^j \tag{4}$$

Proof. We first show that $Z' = \{a \in A \backslash X \mid \exists x \in X, x \succ_{\mathsf{tc}(v)} a\} = Z$. If $a \in A \backslash X$ does not belong to Y then it must be comparable to at least one element in $x \in X$ otherwise we can add it to Y and obtain a larger anti-chain. Hence, since a is not in Y, then $x \succ_{\mathsf{tc}(v)} a$. Also, note that if $a \in Y$ then $a \notin Z'$. This is because if a belonged to both Y and Z, then there exists $x_1, x_2 \in X$ such that $x_1 \succ_{\mathsf{tc}(v)} a$ and $a \succ_{\mathsf{tc}(v)} x_2$ this would mean $x_1 \succ_{\mathsf{tc}(v)} x_2$ which contradicts the anti-chain property of X. For a particular item in X, items in Y are either incomparable to it or must be preferred to it, similarly items in Z are either incomparable or must be dis-preferred to it.

This also implies no item in Z can be preferred over items in Y since if this were to happen, i.e. if $z \succ_{\mathsf{tc}(v)} y$ where $z \in Z, y \in Y$, then $\exists x \in X$ such that $y \succ_{\mathsf{tc}(v)} x$, this implies $z \succ_{\mathsf{tc}(v)} x$ which is impossible from the above observation that $Z \cap Y = \emptyset$.

Consider all rankings $\widetilde{\Omega}(v)$ where we place items of Y at the top, X in the middle and Z at the bottom. Within Y and Z we rank items respecting $\mathsf{tc}(v)$ and since X is an anti-chain, rank these items without restrictions. That is

$$\widetilde{\Omega}(v) = \{r | \forall y \in Y, x \in X, z \in Z, y \succ_r x, x \succ_r z, r|_Y \in \Omega(\mathsf{tc}(v)|_Y), r|_Z \in \Omega(\mathsf{tc}(v)|_Z)\}.$$

Now we argue $\widetilde{\Omega}(v) \subseteq \Omega(v)$. Note that we satisfy preference constraints when ranking within Y, X and Z. Also as we showed above, items in Y are never dis-preferred to items in X or Z and items in X are never dis-preferred to items in Z.

For the lower bound, first observe if $r \in \widetilde{\Omega}(v)$ then $d(r,\sigma) = d(r|_Y, \sigma|_Y) + d(r|_X, \sigma|_X) + d(r|_Z, \sigma|_Z) + \delta$ where δ is defined in the theorem as the number of misorderings of items across X, Y, Z, which is independent of r. Hence,

$$\sum_{r \in \Omega(v)} \phi^{d(r,\sigma)} \geq \sum_{r \in \widetilde{\Omega}(v)} \phi^{d(r,\sigma)} = \phi^{\delta} \left[\sum_{r \in \widetilde{\Omega}(v)} \phi^{d(r|_Y, \sigma|_Y)} \right] \left[\sum_{r \in \widetilde{\Omega}(v)} \phi^{d(r|_X, \sigma|_X)} \right] \left[\sum_{r \in \widetilde{\Omega}(v)} \phi^{d(r|_Z, \sigma|_Z)} \right],$$

Finally, it can be seen that the sum inside the third factor is exactly the normalization constant of an unconstrained Mallows model with $|X| = \mathsf{w}(v)$ items, and hence equal to $\prod_{i=1}^{\mathsf{w}(v)} \sum_{j=0}^{i-1} \phi^j$, the second and fourth factors involve sums over rankings of Y and Z consistent with $\mathsf{tc}(v)$. This proves the lower bound. □

While the lower bound is not in "closed-form" it is useful if $\mathsf{w}(v)$ is large, in other words if there are a sparse number of preference constraints in v (e.g. involving only a small subset of items) we expect $\Omega(v)$ to be large and hence higher probability mass. We fully recover the Mallows model normalization constant if $v = \emptyset$ since $\mathsf{w}(v) = m$. If v is highly constrained—$\Omega(v)$ has smaller probability mass—then $\mathsf{w}(v)$ will be small, but so are the factors involving summations. Note that ϕ^{δ} will decrease whenever there are more comparisons in v that disagree with σ this again corresponds to intuition in the upper bound case.

Corollary 7. *Let L and U be the lower and upper bound as in Theorem 6 and 4, respectively. Then for $r \in \Omega(v)$,*

$$\frac{L}{\prod_{i=1}^m \sum_{j=l_i}^{h_i} \phi^{i-j}} \leq \frac{P(r|v,\sigma,\phi)}{\hat{P}_v(r)} \leq \frac{U}{\prod_{i=1}^m \sum_{j=l_i}^{h_i} \phi^{i-j}} \tag{5}$$

Proof. $\hat{P}_v(r)$ has the form given in Proposition 8 while $P(r|v,\sigma,\phi) \propto \phi^{d(r,\sigma)}$. Then apply upper and lower bounds on the normalizing constant of $P(r|v,\sigma,\phi)$. □

2 Computing a Local Kemenization

Alg. 1 works by first initializing the new σ_k to σ_k^{old} of the previous EM iteration. Then focusing on items x from the top to the bottom of the ranking, successively make adjacent swaps between x and item y above it, whenever the majority of rankings in S_k prefer x over y, otherwise stop swapping and move onto the next item x. This gives a locally optimal ranking: when we finish swapping item x upwards, either x is at the very top or some y is preferred to x by the majority of S_k. In the final constructed ranking y may still be above x in which case x cannot be moved up, if a different y' is above x, then y' must be below x in the initial ranking and was swapped above x because y' is preferred to x by majority in S_k. Hence making an adjacent upward swap for x cannot improve the Kemeny cost. Note that instead of storing all rankings of S_k all we need is its *pairwise tournament graph*: which is a complete directed graph where vertices are A and the weight of each edge $x \to y$, is $c_{xy} = |\{\rho \in S_k : y \succ_\rho x\}|$. This is the "Kemeny cost" of deciding to place x above y.

Algorithm 1 LocalKemeny

Input: $S_k = (\rho_{k1}, \dots, \rho_{kj_k})$

$\sigma \leftarrow \sigma_k^{\text{old}}$
Compute pairwise tournament graph:
for all pair $(x, y) : x, y \in A$ and $x \neq y$ **do**
  $c_{xy} = |\{\rho \in S_k : y \succ_\rho x\}|$.
end for
$d \leftarrow \sum_{\{x,y\} \,:\, x \succ_{\sigma_k} y} c_{xy}$
for $i = 2..m$ **do**
  $x \leftarrow$ item in i-th rank of σ
  for $j = i - 1..1$ **do**
    $y \leftarrow$ item in j-th rank of σ
    if $c_{xy} < c_{yx}$ **then**
      Swap x with y
      $d \leftarrow d - c_{xy} + c_{yx}$
    else
      quit this loop
    end if
  end for
end for

Output: σ, Kemeny cost d

3 Derivation for Non-Parameteric Estimators

This section gives the full derivations of using importance sampling for non-parametric estimators on paired comparison data.

Define a joint distribution q_ℓ over the probability space $\Omega(v_\ell) \times \Omega$,

$$q_\ell(s, r) = \frac{\phi^{d(r,s)}}{|\Omega(v_\ell)| Z_\phi} \tag{6}$$

where Z_ϕ is the Mallows normalization constant with respect to dispersion ϕ. This distribution corresponds to drawing a ranking s uniformly at random from $\Omega(v_\ell)$ and then drawing r according to Mallows with reference ranking s and dispersion ϕ. The estimator, extended to any set of paired comparisons, is

$$\begin{aligned} p(v) &= \frac{1}{n} \sum_{\ell \in N} q_\ell(s \in \Omega(v_\ell), r \in \Omega(v)) \\ &= \frac{1}{n} \sum_{\ell \in N} \sum_{s \in \Omega(v_\ell)} \sum_{r \in \Omega(v)} \frac{\phi^{d(r,s)}}{|\Omega(v_\ell)| Z_\phi}. \end{aligned} \tag{7}$$

Note that this is a distribution over rankings and not incomplete preferences, the above is simply a marginal over $\Omega(v)$. A special case arises when V consists entirely of full rankings, which simplifies to a mixture of Mallows with n equally weighted components each with v_ℓ's as centres and dispersion ϕ. This estimator can be useful for making inferences over the posterior $p(r|v) = p(r)\mathbf{1}[r \in \Omega(v)]/p(v)$ for $r \in \Omega(v)$. Fix v, let $f(s) = \sum_{r \in \Omega(v)} \phi^{d(r,s)}$. Then

$$\begin{aligned} p(v) &= \frac{1}{n Z_\phi} \sum_{\ell \in N} \sum_{s \in \Omega(v_\ell)} \frac{1}{|\Omega(v_\ell)|} f(s) \\ &= \frac{1}{n Z_\phi} \sum_{\ell \in N} \mathop{\mathbb{E}}_{s \sim \Omega(v_\ell)} f(s) \end{aligned}$$

where s is drawn uniformly from $\Omega(v_\ell)$. One can estimate the expectation by importance sampling. Suppose we draw, for each ℓ, rankings $s_\ell^{(1)}, \ldots, s_\ell^{(T)}$ from $\mathsf{AMP}(v_\ell, \sigma, \phi = 1)$ so as to approximate uniform sampling (choose a σ say from $\Omega(v_\ell)$). Let $w_{\ell t} = 1/\hat{P}_v(s_\ell^{(t)})$ Then the estimate is

$$\hat{p}(v) = \frac{1}{nZ_\phi} \sum_{\ell \in N} \frac{\sum_{t=1}^T w_{\ell t} f(s_\ell^{(t)})}{\sum_{t=1}^T w_{\ell t}}.$$

Evaluating $f(s_\ell^{(t)})$ is intractable but can be approximated using our earlier techniques for approximating the log likelihood. In summary we use a nested sampling procedure to first approximate the outer expectation over s and the inner summation $f(s)$.

4 Experiments

We performed five sets of experiments. The first compares how good the posterior sampling method AMP, based on the generalized repeated insertion method, approximates the true Mallows posterior. It turns out to be an excellent approximation. The second experiment compares how good our Monte Carlo evaluation of the log likelihood is. Again, it turns out to be a very good approximation. Building on these two positive results, the last three experiments test our EM algorithm on synthetic data, sushi data, and Movielens data (large m). The synthetic data experiments confirm the effectiveness of our EM algorithm and also reveals insights on how size of preference data (either n or α) impacts learning, and its connections to wisdom of the crowds. Experiments on sushi and Movielens datasets reveal interesting clustering patterns in preferences of agents.

4.1 Approximating Mallows Posterior

For the first set of experiments, we want to get a sense of how well the sampling method AMP approximates the true Mallows posterior $P(r|v, \sigma, \phi)$. In particular, we like to measure the KL divergence of the true posterior to $\hat{P}_v(r)$ which is the distribution defined by the algorithm AMP. We experimented with varying three parameters: number of items m, dispersion parameter ϕ and the fraction of paired comparisons contained in v. The results are show in Fig. 1. We normalized the KL divergence by the entropy of the true Mallows posterior because, for example when increasing m, KL and entropy would corresponding scale. For each setting of the parameters, we generated 20 instances of v according to our probabilistic model, and then evaluated the *exact* KL divergence of the true posterior to $\hat{P}_v$, normalized by the entropy of true posterior. We choose a canonical $\sigma = 12 \cdots m$. The results clearly demonstrate that $\hat{P}_v$ is a very good approximation to the true posterior.

4.2 Evaluating Log Likelihood

We showed the #P-hardness of evaluating the log-likelihood and derived a Monte Carlo estimate based on sampling from AMP. We experimented with how good of an approximation the estimator is. We varied three parameters: (1) number of items m, (2) number of components K, and (3) number of samples T per agent and per component. The results are shown in Fig. 2. In all experiments we fixed number of agents (i.e. number of input preferences) to $n = 50$. For (1), we generated v from a mixture model with $K = 3$, $\boldsymbol{\pi} = (1/3, 1/3, 1/3)$, $\boldsymbol{\sigma}$ drawn uniformly at random K independent times, $\boldsymbol{\phi} = (1/2, 1/2, 1/2)$ and $\alpha = 0.2$. For (2), we generated v from a model with $m = 8$, $\boldsymbol{\pi} = (1/K, \ldots, 1/K)$, $\boldsymbol{\phi} = (1/2, \ldots, 1/2)$, $\boldsymbol{\sigma}$ draw uniformly at random K times and $\alpha = 0.2$. For (3), mixture parameters were $K = 1$, $m = 8$, σ chosen uniformly at random, $\phi = 0.5$ and $\alpha = 0.2$. The parameters for which we evaluated the log likelihood on is generated as follows: $\boldsymbol{\pi}$ sampled from a Dirichlet distribution with parameter being a vector of K 5's. Reference rankings $\boldsymbol{\sigma}$ were drawn uniformly at random, and ϕ drawn uniformly at random in interval $(0, 1)$. Overall the results show that the Monte Carlo approximation is very good, and improves significantly while

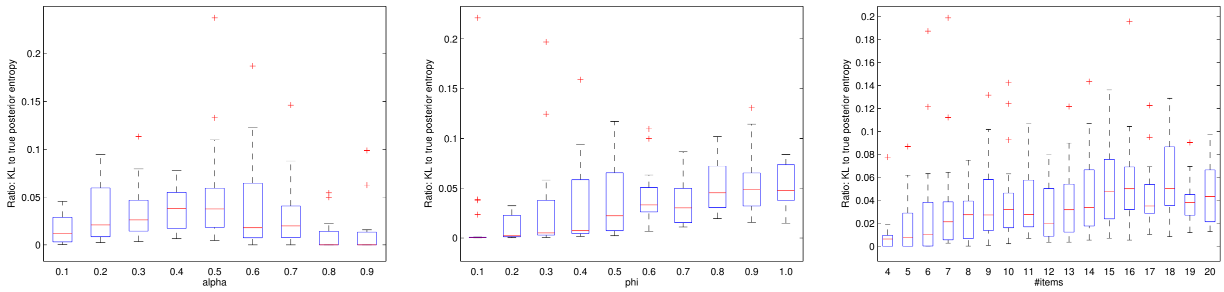

Figure 1: Comparing AMP to the true Mallows posterior. Box and whisker plots with box giving 25-75 percentile of 20 runs, line inside box indicate median and '+' outliers. (1) Varying α, fixing $\phi = 0.5$, $m = 10$ (2) varying ϕ, fixing $\alpha = 0.2$, $m = 10$ (3) Varying m, fixing $\phi = 0.5$ and for $m \leq 13$, $\alpha = 0.2$, for $m > 13$, $\alpha = 0.5$.

reducing variance if we increase the sample size for each agent's log likelihood (as captured by $K \cdot T$), also increasing m slightly degrades the approximation, although it is still an excellent estimate.

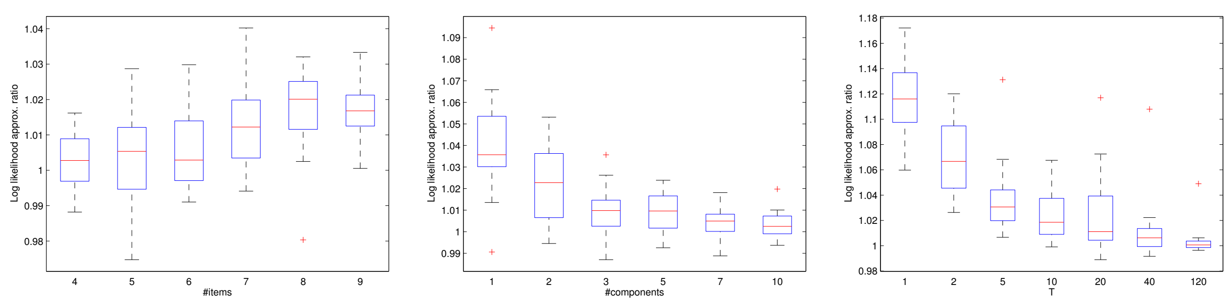

Figure 2: Comparing ratio of true to approximated log likelihood. We ran 20 instances for each setting of parameters. In plot (1) we varied m, fixed $T = 5$ (2) varied K, fixed $T = 5$ (3) varied T.

4.3 EM on Synthetic Data

Having empirically established that sampling procedure AMP is a good approximation to the true posterior, and that the log likelihood can be closely approximated by importance sampling, we can now evaluate how effective our EM algorithm is at recovering parameters in a controlled setting. In our setup, we performed four experiments where we: (1) varied α (2) varied number of items m (3) varied number of components K and (4) varied training data size, that is, number of agent preferences n. For each experiment, we generated random model parameters $\boldsymbol{\pi}$ from a Dirichlet with vector of K 5's, $\boldsymbol{\sigma}$ uniformly at random, and $\boldsymbol{\phi}$ values uniformly at random in the interval $[0.2, 0.8]$. The training data is generated from our probabilistic model using these generated parameters. While varying one parameter for each experiment, we fix the other three, and in particular when fixing the parameters they were always $\alpha = 0.2$, $m = 20$, $K = 3$ and $n = 50K$. Results are shown in Fig. 3. We analyze the performance of EM by (approximately) evaluating the ratio of the log likelihood of the true model parameters $\boldsymbol{\pi}$, $\boldsymbol{\sigma}$, $\boldsymbol{\phi}$ to the EM learned parameters, on test data (preferences) generated from the true model parameters where we chose $n_{test} = n$ and $\alpha_{test} = 1$.

The results suggest that: learning is better when α or n is larger, in other words, when we have more preference data; learning is relatively worse when increasing number of components—because there is less data; and learning improves when increasing m while fixing α—because the transitive closure for larger m gives more preference information (e.g. $a_1 \succ a_2 \succ a_3 \succ a_4 \succ a_5 \succ a_6$ has 5 paired comparisons and is 1/9 of all paired comparisons on $m = 10$ while $a_1 \succ a_2 \succ \cdots \succ a_{100}$ has 99 paired comparisons which is 1/50 of all paired comparisons but its transitive closure is a full ranking).

An interesting implication is the wisdom of the crowds' effect, for example when estimating an objective ranking, i.e. $K = 1$. The amount of data needed for estimating an objective ranking can be traded off by either increasing α, the average number of paired comparisons revealed per agent, or by increasing number of agents n and decreasing α. That is, asking more agents about their objective assessments while decreasing questions per agent, provides roughly the same data needed to find an objective ranking as asking less agents but demanding more objective assessments per agent.

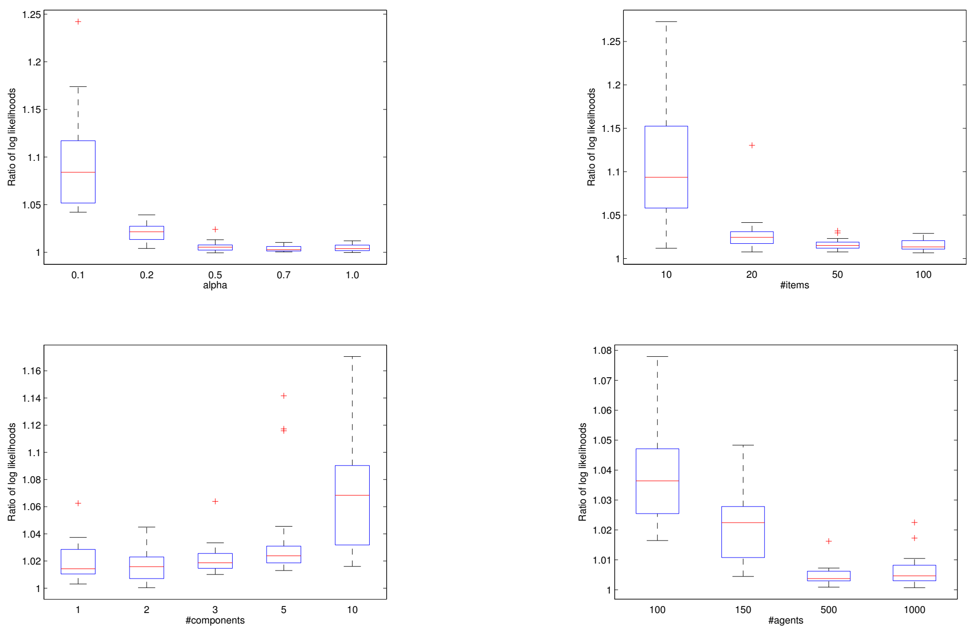

Figure 3: Performance of EM on synthetic dataset. Each plot illustrates the ratio of the log likelihood of true model parameters $\boldsymbol{\pi}$,$\boldsymbol{\sigma}$,$\boldsymbol{\phi}$ to the learned parameters. We ran 20 instances for each setting of experimental parameters. Log likelihoods were approximated by importance sampling with $T = 10$.

4.4 Sushi Data

This dataset consists of sushi preferences surveyed across Japan. We used the first part of the dataset consisting of 5000 complete preferences over $m = 10$ sushi varieties. We split this into 3500 preferences for training and 1500 for validation. Because the full preferences are available, we ran several experiments where we generated training preferences by revealing paired comparisons with probability α. To avoid local maxima, we ran EM ten times (more than what is necessary) for each instance. Fig. 4 shows the results. The plot shows that, even without full preferences, learning is still quite good with only .5 or .4 fraction of all paired comparisons. Learning degrades as less paired preference data becomes available (e.g. going from $\alpha = .3$ to .2). However, there is still enough data for learning with $K = 1, 2$. From the plot it appears $K = 6$ is a good model fit when training on full preferences, Table 1 shows the learned clusters. The pattern emerging is that, with exception of one group, fatty tuna is very well liked. Salmon roe and sea urchin are

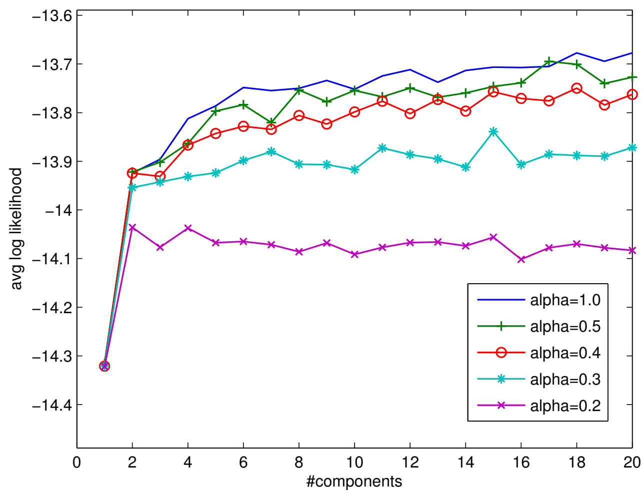

Figure 4: Sushi dataset. Plots of average validation log likelihood when the training data, paired comparisons, are revelead with probabilities $\alpha = .2, .3, .4, .5$. Learning degrades as α gets closer to 0.2, that is, as more paired comparisons are censored.

$\pi_0 = 0.17$ $\phi_0 = 0.66$	$\pi_1 = 0.15$ $\phi_1 = 0.74$	$\pi_2 = 0.17$ $\phi_2 = 0.61$	$\pi_3 = 0.18$ $\phi_3 = 0.64$	$\pi_4 = 0.16$ $\phi_4 = 0.61$	$\pi_5 = 0.18$ $\phi_5 = 0.62$
fatty tuna	shrimp	sea urchin	fatty tuna	fatty tuna	fatty tuna
salmon roe	sea eel	fatty tuna	tuna	sea urchin	sea urchin
tuna	squid	sea eel	shrimp	tuna	salmon roe
sea eel	egg	salmon roe	tuna roll	salmon roe	shrimp
tuna roll	fatty tuna	shrimp	squid	sea eel	tuna
shrimp	tuna	tuna	sea eel	tuna roll	squid
egg	tuna roll	squid	egg	shrimp	tuna roll
squid	cucumber roll	tuna roll	cucumber roll	squid	sea eel
cucumber roll	salmon roe	egg	salmon roe	egg	egg
sea urchin	sea urchin	cucumber roll	sea urchin	cucumber roll	cucumber roll

Table 1: Learned model for $K = 6$ on the sushi dataset with full rankings.

either really liked or hated together, likely because they are not typical "fish meat." Cucumber roll is mostly dispreferred.

4.5 Movielens Data

We applied our EM algorithm on the Movielens dataset[1] to find "preference types" of users. The dataset consists of ∼1 million movie ratings in year 2000 of ∼3900 movies made by ∼6000 users. The ratings were integers from 1 to 5. In our experiments, we focused on the 200 most rated movies. We converted user ratings into paired comparisons as follows: movie $x_1 \succ x_2$ was added to a user's v_ℓ if the rating of x_1 was strictly greater than that of x_2. If the ratings are tied, then they are incomparable and the paired comparison is not added. For example, if A and B had rating 5, C had rating 3 and D rating 1 then the user preference becomes $v = \{A \succ C, A \succ D, B \succ C, B \succ D, C \succ D\}$. We discarded preferences that became empty on the top 200 movies, and used 3986 preferences for training and set aside 1994 for validation. The average number of paired comparisons per user (both training and validation) was roughly 1300.

We ran EM for each component size $K \in \{1, \ldots, 20\}$, and for each K we reran EM 20 times to avoid

[1] see `www.grouplens.org`

local maxima, which is a lot more runs than is needed to avoid local maxima. Then for each K, we took the best run in the sense that the training log likelihood was largest and evaluated the average log likelihood on a validation set whose purpose is in selecting a good K. The log likelihoods were approximated using our Monte Carlo estimate (with $K \cdot T = 120$, i.e. sample size per preference is 120). A C++ implementation was quite fast and resulted in EM running times between 15 to 20 minutes, depending on K (Intel Xeon dual-core 3GHz). The log likelihood plot is shown in Fig. 5. On validation data, the best component sizes were 10 and 5 (with 10 slightly beating 5). While there are various ways to choose the right K (e.g. Bayesian or Akaike information criteria), we use Occam's principle and display the learned components for $K = 5$ in Table 2. This table shows the top 20 movies of each cluster centre as well as the mixture proportions and dispersion paramters.

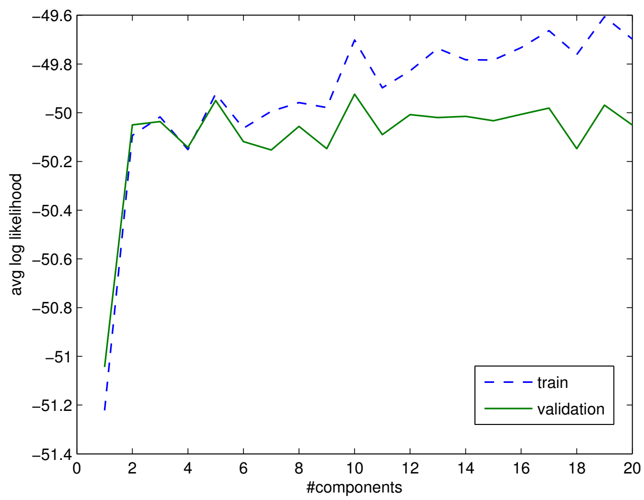

Figure 5: Movielens dataset: average training and validation log likelihoods on the learned model parameters of different component sizes.

References

[1] Graham Brightwell and Peter Winkler. Counting linear extensions is #p-complete. In *ACM Symposium on the Theory of Computing*, pages 175–181, 1991.

[2] Guy Lebanon and Yi Mao. Non-parametric modeling of partially ranked data. *Journal of Machine Learning Research*, 9:2401–2429, 2008.

$\pi_1 = 0.24$, $\phi_1 = 0.98$	$\pi_2 = 0.23$, $\phi_2 = 0.98$	$\pi_3 = 0.21$, $\phi_3 = 0.98$	$\pi_4 = 0.19$, $\phi_4 = 0.98$	$\pi_5 = 0.13$, $\phi_5 = 0.97$
Citizen Kane (1941)	Godfather, The (1972)	Raiders of the Lost Ark (1981)	Shawshank Redemption, The (1994)	Usual Suspects, The (1995)
Godfather, The (1972)	Dr. Strangelove (1963)	Godfather, The (1972)	Life Is Beautiful (1997)	Shawshank Redemption, The (1994)
Dr. Strangelove (1963)	Citizen Kane (1941)	Schindler's List (1993)	Raiders of the Lost Ark (1981)	Schindler's List (1993)
Schindler's List (1993)	Casablanca (1942)	Rear Window (1954)	Schindler's List (1993)	Life Is Beautiful (1997)
Rear Window (1954)	Star Wars: Episode IV - A New Hope (1977)	Star Wars: Episode IV - A New Hope (1977)	Star Wars: Episode IV - A New Hope (1977)	Christmas Story, A (1983)
Shawshank Redemption, The (1994)	Usual Suspects, The (1995)	Shawshank Redemption, The (1994)	Matrix, The (1999)	This Is Spinal Tap (1984)
American Beauty (1999)	Raiders of the Lost Ark (1981)	Casablanca (1942)	Sixth Sense, The (1999)	American Beauty (1999)
Godfather: Part II, The (1974)	Monty Python and the Holy Grail (1974)	Sixth Sense, The (1999)	Sting, The (1973)	Sixth Sense, The (1999)
One Flew Over the Cuckoo's Nest (1975)	Rear Window (1954)	Psycho (1960)	Forrest Gump (1994)	Pulp Fiction (1994)
Casablanca (1942)	Maltese Falcon, The (1941)	Citizen Kane (1941)	Usual Suspects, The (1995)	Princess Bride, The (1987)
Usual Suspects, The (1995)	Blade Runner (1982)	Sting, The (1973)	Braveheart (1995)	Silence of the Lambs, The (1991)
Pulp Fiction (1994)	One Flew Over the Cuckoo's Nest (1975)	Usual Suspects, The (1995)	Green Mile, The (1999)	Godfather, The (1972)
Monty Python and the Holy Grail (1974)	Clockwork Orange, A (1971)	Saving Private Ryan (1998)	Indiana Jones and the Last Crusade (1989)	Forrest Gump (1994)
Fargo (1996)	2001: A Space Odyssey (1968)	Godfather: Part II, The (1974)	Saving Private Ryan (1998)	Fight Club (1999)
Life Is Beautiful (1997)	North by Northwest (1959)	Silence of the Lambs, The (1991)	Princess Bride, The (1987)	Fargo (1996)
Graduate, The (1967)	Pulp Fiction (1994)	Wizard of Oz, The (1939)	Star Wars: Episode V - The Empire Strikes Back (1980)	Ferris Bueller's Day Off (1986)
North by Northwest (1959)	Godfather: Part II, The (1974)	Dr. Strangelove (1963)	Silence of the Lambs, The (1991)	Raising Arizona (1987)
GoodFellas (1990)	Chinatown (1974)	Jaws (1975)	Good Will Hunting (1997)	Saving Private Ryan (1998)
Chinatown (1974)	Apocalypse Now (1979)	Braveheart (1995)	Ferris Bueller's Day Off (1986)	Good Will Hunting (1997)
Raiders of the Lost Ark (1981)	Shawshank Redemption, The (1994)	Aliens (1986)	When Harry Met Sally (1989)	Matrix, The (1999)

Table 2: Learned model for $K = 5$ on Movielens. Shows the top 20 (out of 200) movies.