Active embedding search via noisy paired comparisons

Gregory H. Canal, Andrew K. Massimino, Mark A. Davenport, and Christopher J. Rozell
Georgia Institute of Technology

I. Introduction

We consider the task of user preference learning, where we are given a set of *items* embedded in a low-dimensional Euclidean space and aim to represent the preferences of a *user* as a continuous point in the same space so that their preference point is close to items the user likes and far from items the user dislikes. The recovered preference point can be used in various tasks, for instance in the recommendation of nearby items, clustering of users with similar preferences, or personalized product creation. To estimate this point, we consider a system using the *method of paired comparisons*, where during a sequence of interactions a user chooses which of two presented items they prefer [1]. In this work, we assume a response model common in psychometrics literature [2], where the probability of a user located at w choosing item p over item q in a paired comparison is given by

$$P(p \prec q) = f(k_{pq}(a^T w - b)), \tag{1}$$

where $a = 2(p-q)$ and $b = \|p\|^2 - \|q\|^2$ encode the normal vector and threshold of a hyperplane bisecting items p and q, $f(x) = 1/(1+e^{-x})$ is the logistic function, and k_{pq} is the pair's *noise constant*, which represents the signal-to-noise ratio of a particular query. Querying all possible pairs to estimate user preferences is not only prohibitively expensive for large datasets, but also unnecessary since not all queries are informative. The main contribution of this work is the design and analysis of two new query selection algorithms for low-dimensional pairwise search that select the most informative pairs by directly modeling redundancy and noise in user responses. While previous active algorithms exist for related paired comparison models [3], [4], none directly account for probabilistic user behavior as we do here.

II. Query Selection

For the i^{th} paired comparison involving items $p_i, q_i \in \mathbb{R}^d$ $(d \geq 2)$, let $Y_i = 1$ (resp. 0) denote a preference for p_i (resp. q_i). After i queries, we have the vector of responses $Y^i = \{Y_1, Y_2, \dots Y_i\}$, with each response assumed to be conditionally independent from previous responses when conditioned on preference $W \in \mathbb{R}^d$, which is assumed to be drawn from a uniform hypercube prior. Let $\Sigma_{W|Y^i} \equiv \mathbb{E}[(W - \mathbb{E}[W|Y^i])(W - \mathbb{E}[W|Y^i])^T|Y^i]$, and define the *posterior volume* as $|\Sigma_{W|Y^i}|$. We aim to adaptively select queries based on previous responses that minimize the mean-squared error (MSE) of a Bayesian preference estimator. Although selecting queries to directly minimize MSE is computationally expensive, under the model in (1) it can be shown that low differential entropy of the preference posterior is a necessary and sufficient condition for low posterior volume, which itself is a necessary condition for low MSE. This suggests a strategy of selecting queries that maximize the decrease in posterior entropy after a query (referred to here as the *information gain* and denoted by $I(W; Y_i|y^{i-1})$) [5]. Based on this notion, we develop two strategies that mimic the action of maximizing information gain while being analytically and computationally tractable, respectively.

Consider the i^{th} selected pair with bisecting hyperplane parameterized by (a_i, b_i) and define an *equiprobable* query strategy to select b_i such that each item in the query will be chosen by the user with probability $\frac{1}{2}$, and a *mean-cut* strategy to select b_i such that the query hyperplane passes through the posterior mean. Define a query's *projected variance* as the variance of the posterior marginal in the direction of a query's hyperplane, i.e. $a_i^T \Sigma_{W|y^{i-1}} a_i$. With these definitions, we have the following result:

Proposition II.1. *For both equiprobable and mean-cut queries, information gain is nearly-tightly lower bounded by monotonically increasing functions of projected variance.*

This result suggests choosing a_i which *maximize projected variance*. We refer to the selection of equiprobable queries in the direction of largest projected variance as the *equiprobable-max-variance* (EPMV) scheme, and mean-cut queries in the direction of largest projected variance as the *mean-cut, maximum variance* (MCMV) scheme. Our primary result concerns the expected number of comparisons sufficient to reduce the posterior volume below a specified threshold set a priori, using EPMV.

Theorem II.2. *For the EPMV query scheme with each selected query satisfying $k_i\|a_i\| \geq k_{min}$ for some constant $k_{min} > 0$, consider the stopping time $T_\varepsilon = \min\{i : |\Sigma_{W|y^i}|^{\frac{1}{d}} < \varepsilon\}$ for stopping threshold $\varepsilon > 0$. We have*

$$\mathbb{E}[T_\varepsilon] = O\left(d \log \frac{1}{\varepsilon} + \left(\frac{1}{\varepsilon k_{min}^2}\right) d^2 \log \frac{1}{\varepsilon}\right).$$

Furthermore, for any query scheme, $\mathbb{E}[T_\varepsilon] = \Omega\left(d \log \frac{1}{\varepsilon}\right)$.

This result has a favorable dependence on the dimension d, and the upper bound can be interpreted as a blend between two rates, one that matches a generic lower bound and another that decreases to zero for large noise constants. On the other hand, MCMV is attractive from a computational standpoint since the posterior mean and covariance can be estimated *once* per query round, and subsequent calculation of the hyperplane distance from mean and projected variance requires only $O(d^2)$ computations per pair, which scales more favorably than both the information gain maximization (*InfoGain*) and EPMV strategies.

III. Results

To evaluate our methods, we constructed a realistic embedding consisting of multidimensional item points from the Yummly `Food-10k` dataset of [6], [7], and simulated preference search over randomly generated preference points and user responses, as depicted in Figure 1. We compare against two competing methods that we refer to as GaussCloud-Q [3] and ActRank-Q [4], as well as against randomly selected queries. We evaluate each method on both the logistic noise model ("matched" noise) as well as a scenario where noise is generated according to a Gaussian model while the Bayesian methods continue to calculate the posterior as if the responses were logistic ("mismatched" noise). Our experiments demonstrate that both InfoGain approximation strategies (EPMV and MCMV) significantly outperform the state-of-the-art methods in active preference estimation in the context of low-dimensional item embeddings with noisy user responses, and perform similarity to InfoGain, the technique they were designed to approximate. This is true even when generating responses according to a different model than the one used for Bayesian estimation.

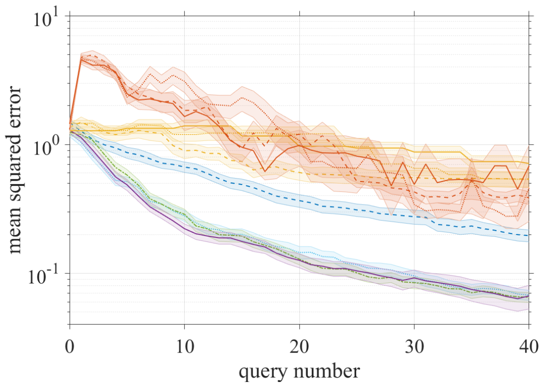

(a) Estimation error: matched logistic noise, $d = 4$

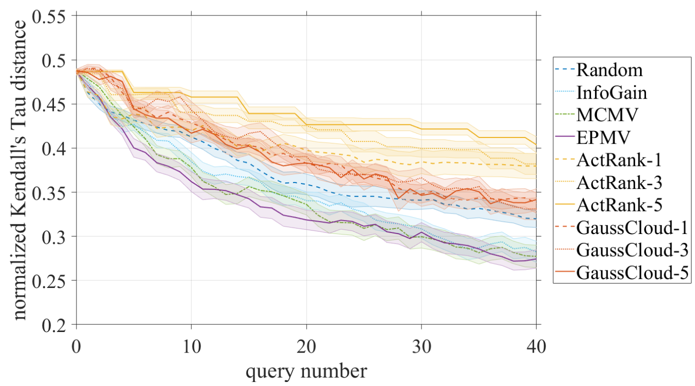

(b) Ranking performance: matched logistic noise, $d = 4$

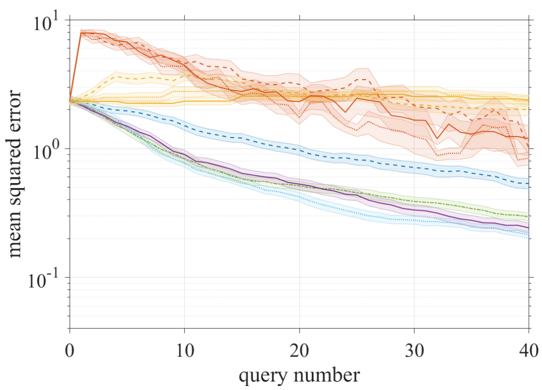

(c) Estimation error: mismatched Gaussian noise, $d = 7$

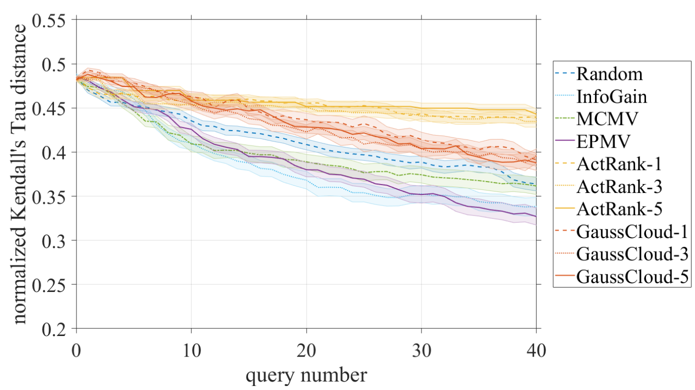

(d) Ranking performance: mismatched Gaussian noise, $d = 7$

Fig. 1: Performance metrics in evaluating preference searching over 40 queries, averaged over 50 trials per method in a search task of the Yummly `Food-10k` dataset. For Random, InfoGain, MCMV, and EPMV, we estimate the user point as the posterior mean since this is the MMSE estimator. All traces are plotted with $\pm$ one standard error. (Left Column) MSE in estimating user preference. (Right Column) for each trial, a batch of 15 items was uniformly sampled without replacement from the dataset, and the normalized Kendall's Tau distance was calculated between a ranking of these items by distance to the ground truth preference point and a ranking of distance to the estimated point, with lower distance indicating better performance. To get an unbiased estimate, this metric is averaged over 1000 batches per trial, and error bars calculated with respect to the number of trials. Rather than solely measuring preference estimation error, this metric measures performance in the context of a recommender system type task, which is a common application of preference learning. (Top Row) "normalized" ($k_{pq} = k_0\|(p-q)\|^{-1}$) logistic model with matching noise in $d = 4$. (Bottom Row) "decaying" ($k_{pq} = k_0 \exp(-2\|(p-q)\|)$) logistic model with mismatched Gaussian "normalized" noise in $d = 7$. For a complete set of results, see the full paper at https://arxiv.org/abs/1905.04363.

Acknowledgments

This work was supported by NSF grants CCF-1350616 and CCF-1350954, ONR grant N00014-15-1-2619, and a gift from the Alfred P. Sloan Foundation.

References

[1] H. A. David, *The method of paired comparisons*. London, 1963, vol. 12.
[2] R. A. Bradley and M. E. Terry, "Rank analysis of incomplete block designs: I. the method of paired comparisons," *Biometrika*, vol. 39, no. 3/4, pp. 324–345, 1952.
[3] A. K. Massimino and M. A. Davenport, "As you like it: Localization via paired comparisons," *arXiv preprint arXiv:1802.10489*, 2018.
[4] K. G. Jamieson and R. Nowak, "Active ranking using pairwise comparisons," in *Advances in Neural Information Processing Systems*, 2011, pp. 2240–2248.
[5] T. M. Cover and J. A. Thomas, *Elements of information theory*. John Wiley & Sons, 2012.
[6] M. Wilber, I. S. Kwak, D. Kriegman, and S. Belongie, "Learning concept embeddings with combined human-machine expertise," in *Proceedings of the IEEE International Conference on Computer Vision*, 2015, pp. 981–989.
[7] M. J. Wilber, I. S. Kwak, and S. J. Belongie, "Cost-effective hits for relative similarity comparisons," in *Second AAAI conference on human computation and crowdsourcing*, 2014.