Real-time Multiattribute Bayesian Preference Elicitation with Pairwise Comparison Queries

Shengbo Guo
ANU/NICTA
Locked Bag 8001
Canberra ACT 2601, Australia
Shengbo.Guo@nicta.com.au

Scott Sanner
NICTA/ANU
Locked Bag 8001
Canberra ACT 2601, Australia
Scott.Sanner@nicta.com.au

1 Introduction

Preference elicitation (PE) is an important component of interactive decision support systems that aim to make optimal recommendations to users by actively querying their preferences. The PE task consists of (a) querying the user about their preferences and (b) recommending an item that maximizes the user's latent utility. Of course, a PE system is limited by real-world performance constraints that require phase (a) to be efficient while ensuring phase (b) can make an optimal recommendation with high certainty. Bayesian approaches to PE [2] have received interest in recent years due to their robust handling of noise in the elicitation process, however, previous work has either relied on expensive sampling methods [2] or on expensive EM refitting of mixture models [1] to deal with the lack of a closed-form for the utility belief update. In this work, we propose to avoid both of these problems by adapting the Bayesian ranking approach of TrueSkill [3] to multiattribute Bayesian PE, which allows us to efficiently maintain and update the belief representation in real-time and naturally facilitates the efficient evaluation of value of information (VOI) heuristics for use in query selection strategies. Our best VOI query strategy is both space- and time-efficient (in contrast to related work) and performs on par with the most accurate (and often computationally intensive) algorithms on experiments with a real-world dataset.

2 Bayesian Preference Elicitation

In multiattribute utility theory (MAUT), utilities are modeled over a D-dimensional *attribute set* $\mathcal{X} = \{X_1, \ldots, X_D\}$ with *attribute choices* $X_d = \{x_{d1}, \ldots, x_{d|X_d|}\}$ (where $|X_d|$ denotes the cardinality of X_d). An *item* is described by its attribute choice assignments $\mathbf{x} = (x_1, \ldots, x_D)$ where $x_d \in X_d$. In our model, an *attribute weight vector* $\mathbf{w} = (w_{11}, \ldots, w_{1|X_1|}, \ldots, w_{D1}, \ldots, w_{D|X_D|})$ describes the utility of *each* attribute choice in *each* attribute dimension. Furthermore, we assume that the utility $u(\mathbf{x}|\mathbf{w})$ of item $\mathbf{x}$ w.r.t. attribute weight vector $\mathbf{w}$ decomposes *additively* over the attribute choices of $\mathbf{x}$, i.e.,

$$u(\mathbf{x}|\mathbf{w}) = \sum_{d=1}^{D} \mathbf{w}_{d,\#(\mathbf{x},d)}, \quad u^*(\mathbf{x}) = \sum_{d=1}^{D} \mathbf{w}^*_{d,\#(\mathbf{x},d)} \tag{1}$$

where $\#(\mathbf{x}, d)$ returns index in $\{1, \ldots, |X_d|\}$ for attribute choice x_d of $\mathbf{x}$ and u^* represents the user's true utility w.r.t. their true (but hidden) $\mathbf{w}^*$. We take a Bayesian perspective on learning $\mathbf{w}$ and thus maintain a probability distribution $P(\mathbf{w})$ representing our beliefs over $\mathbf{w}^*$.

Because $P(\mathbf{w})$ is a distribution over a multidimensional continuous random variable $\mathbf{w}$, we represent this distribution as a Gaussian with diagonal covariance, represented compactly in a factorized format as follows:

$$P(\mathbf{w}) = \prod_{d=1}^{D} \prod_{i=1}^{|X_d|} p(w_{di}) = \prod_{d=1}^{D} \prod_{i=1}^{|X_d|} \mathcal{N}(w_{di}; \mu_{di}, \sigma^2_{di}). \tag{2}$$

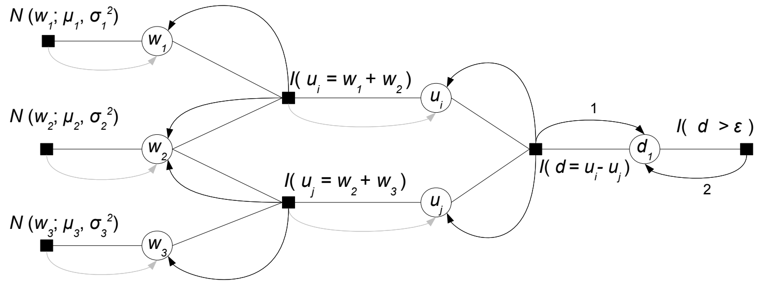

Figure 1: PE factor graph variant of TrueSkill for $q_{ij} = i \succ j$. Items i and j have two attribute choices each with respective weights (w_1, w_2) and (w_2, w_3) (note that i and j share the common attribute choice with weight w_2). The posterior over (w_1, w_2, w_3) can be inferred with the following message passing schedule: (1) messages pass along *gray* arrows from left to right, (2) the marginal over d is updated via message 1 followed by message 2 (which required moment matching), (3) messages pass from right to left along *black* arrows.

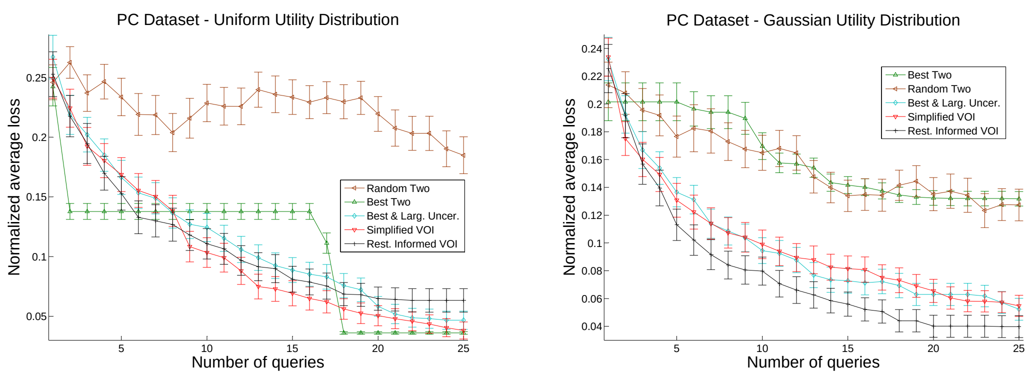

Figure 2: Expected loss vs. number of queries for various PE strategies on the PC dataset. Error bars indicate standard error. Results with uniform utility (left) and diagonal Gaussian utility (right).

We take a Bayesian approach to PE. Thus, given a prior utility belief $P(\mathbf{w}|R^n)$ w.r.t. a (possibly empty) set of $n \geq 0$ query responses $R^n = \{q_{kl}\}$ and a new query response q_{ij}, we perform the following Bayesian update to obtain a posterior belief $P(\mathbf{w}|R^{n+1})$ where $R^{n+1} = R^n \cup \{q_{ij}\}$:

$$\begin{aligned} P(\mathbf{w}|R^{n+1}) &\propto P(q_{ij}|\mathbf{w}, R^n)P(\mathbf{w}|R^n) \\ &\propto P(q_{ij}|\mathbf{w})P(\mathbf{w}|R^n) \end{aligned} \tag{3}$$

We adapt the TrueSkill inference framework [3], and present an approximate message-passing approach to estimate the posterior utility belief given the query result as illustrated in Figure 1.

In Figure 2, we show a plot of the normalized average loss (of recommending the current best item in expectation) of all algorithms vs. the number of query responses elicited on a *PC* dataset consisting 120 items each with 8 attributes. The key observations are that (1) the value of information (VOI) heuristics always perform the best, and (2) in particular, the Restricted Informed VOI has excellent real-time performance (times not shown here due to space limitations).

References

[1] C. Boutilier. A POMDP formulation of preference elicitation problems. In *AAAI*, pages 239–246, 2002.

[2] U. Chajewska, D. Koller, and R. Parr. Making rational decisions using adaptive utility elicitation. In *AAAI*, pages 363–369, 2000.

[3] R. Herbrich, T. Minka, and T. Graepel. TrueSkill^{TM}: A Bayesian skill rating system. In *NIPS*, pages 569–576, 2006.