

---

# Monte Carlo methods for Utility Elicitation and Prior Acquisition

---

**Paolo Viappiani\***

Department of Computer Science  
University of Regina  
paolo.viappiani@gmail.com

**Craig Boutilier**

Department of Computer Science  
University of Toronto  
cebly@cs.toronto.edu

## Abstract

Utility elicitation is an important component of many applications, as decision support systems and recommender systems. Such systems query the users about their preferences and give recommendations based on the system's belief about the utility function.

Critical to these applications is the acquisition of prior distribution about the utility parameters, and the possibility of real time Bayesian inference. In this paper we consider Monte Carlo methods for these problems.

## 1 Bayesian Utility Elicitation

Utility elicitation is a key component in many decision support applications and recommender systems, since appropriate decisions or recommendations depend critically on the preferences of the user on whose behalf decisions are being made. Since full elicitation of user utility is prohibitively expensive in most cases (w.r.t. time, cognitive effort, etc.), we must often rely on partial utility information. This is the case of *interactive preference elicitation*.

As a user's utility function will not be known with certainty, following recent models of Bayesian elicitation [2, 1, 3, 5], the system's knowledge about the user preferences is represented as probabilistic *beliefs*. Interactive elicitation must selectively decide which queries are most informative relative to the goal of making good or optimal recommendations and then, following user responses, update the distribution. An important requirement for utility elicitation is that inference can be made *real-time*, the system needs to output a recommendation or ask a query in no more than a few seconds.

While there are a variety of query types that can be used, comparison queries are especially natural, asking a user if she prefers one option to another. As the number of items in a dataset can be extremely large, iterating over all possible comparison is unfeasible. Recently [5] we showed that, under very general assumptions, the optimal choice query w.r.t. the expected value of information (EVOI) coincides with optimal recommendation set, that is, a set maximizing expected utility of the user selection (a simpler and submodular problem). Based on this, we can provide algorithms that select near-optimal comparison queries with worst-case guarantees; we also considered a local search technique that can select the query to ask in a fraction of a second, even for datasets with several hundreds of items.

These strategies for query optimization (and similar approximated strategies [3]) rely on the assumptions that prior utility information is available and that inference is fast enough so that user answers can be used as additional knowledge for further elicitation. As users are not usually willing to wait more than a couple of seconds, effective inference methods are therefore crucial.

---

\*From 01/2011 onwards at Aalborg University.

Bayesian inference is challenging because common prior distributions are not closed under Bayesian inference for most types of preference queries (in particular for comparison queries). Therefore approximated inference is required, and Monte Carlo methods provide a viable solution.

Another crucial problem for utility elicitation is the acquisition of prior information. The value of Bayesian approaches to utility elicitation is severely hindered if an incorrect prior is used. In this paper we consider how Monte Carlo methods can be used in utility-based recommendation systems with the following purposes 1) to update inference about the user’s possible utility function (given the user’s query responses), and 2) to acquire a prior distribution about utility parameters given preference statements from previous users.

**The Underlying Decision Problem** The system is charged with the task of recommending an option to a user in some multiattribute space, for instance, the space of possible product configurations from some domain (e.g., computers, cars, apartment rental, etc.). Products are characterized by a finite set of attributes  $\mathcal{X} = \{X_1, \dots, X_n\}$ , each with finite domain  $Dom(X_i)$ . For instance, attributes may correspond to the features of various cars, such as color, engine size, fuel economy, etc., with  $\mathbf{X}$  defined either by constraints on attribute combinations. The user has a *utility function*  $u : Dom(\mathcal{X}) \rightarrow \mathbf{R}$ . The precise form of  $u$  is not critical, but we assume that  $u(\mathbf{x}; w)$  is parametric in  $w$ . We often refer to  $w$  as the user’s “utility function” for simplicity, assuming a fixed form for  $u$ . Given a choice set  $S$  with  $x \in S$ , let  $S \triangleright x$  denote that  $x$  has the greatest utility among the items in  $S$  (for a given utility function  $w$ ).

The system’s uncertainty about the user preferences is reflected in a distribution, or *beliefs*,  $P(w; \theta)$  over the space  $W$  of possible utility functions. Here  $\theta$  denotes the parameterization of our model, and we often refer to  $\theta$  as our *belief state*. Given  $P(\cdot; \theta)$ , we define the *expected utility* of an option  $x$  to be  $EU(x; \theta) = \int_W u(x; w)P(w; \theta)dw$ . If required to make a recommendation given belief  $\theta$ , the optimal option  $x^*(\theta)$  is that with greatest expected utility  $EU^*(\theta) = \max_{x \in X} EU(x; \theta)$ , with  $x^*(\theta) = \arg \max_{x \in X} EU(x; \theta)$ . When the user selects an option  $x$  in a choice set  $S$ , the belief is updated to  $P(w; \theta | S \rightsquigarrow x)$ . For query selection strategies, we refer to [5].

**Probabilistic Response Model** In utility elicitation, the user’s response to a choice set tells us something about her preferences; but this depends on the *user response model*. For any choice set  $S$  with  $x_i \in S$ , let  $S \rightsquigarrow x_i$  denote the event of the user selecting  $x_i$ . A response model  $R$  dictates, for any choice set  $S$ , the probability  $P_R(S \rightsquigarrow x_i; w)$  of any selection given utility function  $w$ . We consider three possible response models for choice queries.

In the *noiseless response model*, the user is always able to identify the preferred item in a choice query set; thus  $P_R(S \rightsquigarrow x; 1) = 1$  if  $w$  is such that  $x$  has higher utility than any other in the choice set, 0 otherwise. The set of feasible utility functions is refined by imposing  $k - 1$  linear constraints of the form  $u(x_i; w) \geq u(x_j; w)$ ,  $j \neq i$ , and the new belief state is obtained by restricting  $\theta$  to have non-zero mass only on  $W \cap S \triangleright x_i$  and renormalizing. The *constant noise model* instead assumes each option  $x$ , apart from the most preferred option  $x_w^*$  relative to  $w$ , is selected with (small) constant probability  $P_C(S \rightsquigarrow x; w) = \beta$ , with  $\beta$  independent of  $w$ .

Finally the *logistic response model*  $R_L$  is commonly used in choice modeling, and is variously known as the Luce-Sheppard, Bradley-Terry, or *mixed multinomial logit* model. Selection probabilities are given by  $P_L(S \rightsquigarrow x; w) = \frac{\exp(\gamma u(x; w))}{\sum_{y \in S} \exp(\gamma u(y; w))}$ , where  $\gamma$  is a temperature parameter. For comparison queries (i.e.,  $|S| = 2$ ),  $P_L$  is the logistic function of the difference in utility between the two options.

## 2 Monte Carlo Methods

**Online Inference** During the interaction with the user, the system needs to update the belief taking into account user responses (for instance, the user select  $x$  as the preferred outcome in slate  $S$ ), resulting in a new distribution  $P(w; \theta | S \rightsquigarrow x)$ . Our methods uses particles to represent assignments to the utility parameters  $w$ , initially generated according to the given prior. As we are in online settings, every time the user answers a query we can propagate the particles with importance sampling. Particle weights are determined by applying the response model to observed responses:  $P_R(S \rightsquigarrow x; w)$  (the selection probability) where  $S$  is the choice set and  $\mathbf{x}$  the outcome selected.

To overcome the problem of particle degeneration (most particles eventually have low or no weight), we use slice-sampling [4] to regenerate particles w.r.t. to the response-updated belief state  $\theta$  whenever the *effective number of samples* drops significantly. The choice of the best mixing between importance sampling and slice sampling is an open question, as the number of necessary particles. With 50000 particles in standard elicitation problems, importance sampling requires less than 1 second, but particle regeneration requires around 30 seconds.

In the case of noiseless responses an alternative to regenerate particles is Gibbs-sampling, as the region of feasible utilities can be represented by a convex region. Given dimension  $j$ , we need to sample according to the prior in the interval  $[w_j^\perp, w_j^\top]$  where  $w_j^\perp, w_j^\top$  can be derived analytically.

**Learning Utility Priors from Data** We want to acquire a prior distribution for factored utilities in multi attribute domain, to be used for utility elicitation. We are given as input a number of *preference statements* for several users, of the type  $S_i^j \rightsquigarrow x_i^j$  (answers to preference queries from previous users) where  $S_i^j$  is the  $j$ -th query choice set shown to user  $i$ , and  $x_i^j$  his selection.

As before, we assume a given response model  $R$  for the users, that dictates for any choice set  $S$ , the probability of selection  $P_R(S \rightsquigarrow x_i; w)$ . Our algorithm iterates over all the users and perform importance sampling using the joint likelihood of all preference statements. In the following, a particle is a vector of weights uniquely identifying a utility function ( $n$  is the number of particles,  $m$  the dimensionality of the utility parameters).

**Algorithm:**

1. Sample  $n$  particles uniformly from  $U[0, 1]^m$  and call  $D_0$  this initial particle set
2. For each user  $i$ :
  - *Importance Sampling* step: resample  $n$  particles from  $D_0$ , with weights according to:  $\prod_j P_R(S_i^j \rightsquigarrow x_i^j; w)$
  - Call  $D^i$  the resulting user-specific distribution of particles
3. Merge the particles obtained with all the users:  $D^\top = \bigcup D^i$
4. Cluster  $D^\top$  according to a *mixture gaussian model* (for instance expectation-maximization EM)

The last step (4) is optional and gives us a compact representation of the prior. Another possibility is to repeat several times the steps 2-4, sampling from the output of the previous iterations.

**Learning the Response Model** We are considering how to learn the parameter  $\gamma$  of the logistic response model from . A natural assumption is to consider the value of  $\gamma$  shared from all users. In this case, we can learn utility priors for different values of  $\gamma$  and select the value that maximizes the overall likelihood of the preference statements. An alternative is to augment utility particles with an hypothesis about the  $\gamma$  and resample based on the likelihood of responses given both  $w$  and  $\gamma$ . We are currently experimenting these approaches in a real dataset<sup>1</sup> of preference rankings<sup>2</sup> .

**References**

[1] Craig Boutilier. A POMDP formulation of preference elicitation problems. In *Proceedings of the Eighteenth National Conference on Artificial Intelligence (AAAI-02)*, pages 239–246, Edmonton, 2002.

[2] Urszula Chajewska, Daphne Koller, and Ronald Parr. Making rational decisions using adaptive utility elicitation. In *Proceedings of the Seventeenth National Conference on Artificial Intelligence (AAAI-00)*, pages 363–369, Austin, TX, 2000.

[3] Shengbo Guo and Scott Sanner. Real-time multiattribute bayesian preference elicitation with pairwise comparison queries. In *Proceedings of the 13th International Conference on Artificial Intelligence and Statistics (AISTATS-10)*, Sardinia, Italy, 2010.

[4] Radford M. Neal. Slice sampling. *The Annals of Statistics*, 31(3):705–70, 2003.

[5] Paolo Viappiani and Craig Boutilier. Optimal bayesian recommendation sets and myopically optimal choice query sets. In *Neural Information Processing Systems (NIPS)*. MIT press, 2010.

<sup>1</sup><http://www.kamishima.net/sushi/>

<sup>2</sup>When the input are in the form of rankings  $(x_1, x_2, \dots, x_k)$ , ordered outcomes from most to least preferred, the likelihood (when the user’s utility is  $w$ ) is given by  $\prod_{i,j:j>i} P_R(\{x_i, x_j\} \rightsquigarrow x_i; w)$ .