Recommendation Sets and Choice Queries: There Is No Exploration/Exploitation Tradeoff!

Paolo Viappiani
Department of Computer Science
Aalborg University
Aalborg East, Denmark
paolo@cs.aau.dk

Craig Boutilier
Department of Computer Science
Univeristy of Toronto
Toronto, ON, Canada
cebly@cs.toronto.edu

Abstract

Utility elicitation is an important component of many applications, such as decision support systems and recommender systems. Such systems query users about their preferences and offer recommendations based on the system's belief about the user's utility function. We analyze the connection between the problem of generating optimal recommendation sets and the problem of generating optimal choice queries, in the context of both Bayesian and regret-based elicitation. Our results show that, somewhat surprisingly, under very general circumstances, the optimal recommendation set coincides with the optimal query.[1]

Adaptive Utility Elicitation

Preference elicitation is a challenging task for a number of reasons. First of all, full elicitation of user preferences is prohibitively expensive in most cases (w.r.t. time, cognitive effort, etc.) and we must often rely on partial information. Second, many decision problems have large outcome or decision spaces; techniques for elicitation and recommendation must therefore be scalable. Third, it should be easy for users to provide information about their preferences, possibly accounting for noisy responses.

Adaptive utility elicitation (Braziunas and Boutilier 2008) tackles these challenges by representing the system knowledge about the user in form of *beliefs*, that are updated following user responses. Elicitation queries can be chosen adaptively given the current belief. In this way, one can often make good (or even optimal) recommendations with sparse knowledge of the user's utility function.

There are two main frameworks for representing utility uncertainty. In one approach (Boutilier et al. 2006) the system maintains an explicit representation of a set of *feasible* utility functions, usually represented compactly by constraints; recommendations are generated using the *minimax regret* criterion. Alternatively, a pure Bayesian approach (Chajewska et al. 2000; Boutilier 2002) places a probabilistic prior over the possible utility functions (typically in the form of a density over utility function parameters), and updates the distribution based on observations; the option with greatest expected utility is recommended.

Given the current belief about the user's utility function, it is important to select good queries so that recommendations can quickly improve. In the case of Bayesian preference elicitation, a natural criterion for queries is *expected value of information (EVOI)*; it is, however, extremely expensive computationally and, because of this, most approaches select queries using heuristics with no theoretical guarantees. For regret-based elicitation, we introduce a non-probabilistic analogue of EVOI. Such informative criteria can often generate better queries than strategies that aim at reducing uncertainty *per se*, such as entropy-based methods.

In both frameworks, we show how (myopically) optimal or near-optimal queries can be generated by exploiting the connection with the problem of generating an optimal recommendation set.

An Exploration/Exploitation Tradeoff?

In general, there is a tension between making good recommendations for the user and eliciting "useful" information from the user. Intuitively, in order to make a good recommendation the system should exploit its current knowledge of the utility function. On the other hand, when asking queries, the system aims to acquire more information, so better recommendations can be made in the future.

Since utility is uncertain, there is often value in recommending a *set* of options from which the user can choose her most preferred. Picking a "diverse" set of recommended options increases the odds of recommending at least one item with high utility. Intuitively, such a set of "shortlisted" recommendations should include options that are *diverse* in the following sense: recommended options should be highly preferred relative to a wide range of "likely" user utility functions (relative to the current belief) (Price and Messinger 2005; Boutilier, Zemel, and Marlin 2003). This stands in contrast to some recommender systems that define diversity relative to product attributes (Reilly et al. 2005), with no direct reference to beliefs about user utility. It is not hard to see that "top k" systems, those that present the k options with highest expected utility, do not generally result in good recommendation sets (Price and Messinger 2005).

Resolving this tension is also important because in online recommendation systems, options can be shown with dual

[1] We summarize the key contributions of two earlier papers (Viappiani and Boutilier 2009; 2010), providing a unified view.

goal of recommendation and elicitation. This is common, for instance, in conversational recommender systems such as *Critique Shop* (Reilly et al. 2007), where several products are displayed, and user selections are not only used to update the preference model, but might also lead to purchase decisions. Since a user can end the interaction at any time, it is important that the system always shows products that are likely to be highly preferred.

Among the many possible types of queries, we focus on *choice queries*. Such queries are commonly used in conjoint analysis and product design (Louviere et al. 2000), requiring a user to indicate which choice/product is most preferred from a set of k options. Hence, we can view any set of products as either a recommendation set or query (or choice) set. Given a set, one can ask: what is the value of the set viewed as recommendation set; or what is its value as a query? We consider the connection between these two criteria in both the Bayesian and the regret-base elicitation frameworks. In both cases we show that, quite surprisingly, under very general assumptions, the optimal recommendation set is also an optimal query set.

Formal Model

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. Options or products are characterized by a finite set of attributes $\mathcal{X} = \{X_1, ...X_n\}$, 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.; $\mathbf{X} \subseteq Dom(\mathcal{X})$ is the set of *feasible configurations*, defined 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 $w \in W$. 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). Note that the regions $W \cap S \triangleright x_i, x_i \in S$, partition utility space (ignoring "ties" which are easily dealt with, but complicate presentation). If we assume that the user correctly identifies the preferred item when presented with a choice query (noisy responses are discussed below), the choice of $x_i \in S$ refines the set of feasible utility functions W by imposing $k-1$ constraints of the form $u(x_i; w) \geq u(x_j; w), j \neq i$. We use $S \rightsquigarrow x_i$ to denote the event of the user selecting x_i among the items in S.

The mathematical analysis now diverges depending on whether one assumes a Bayesian (using a density over possible utility parameters) or a non-Bayesian approach based on strict uncertainty. We provide decision and elicitation criteria for both cases (summarized in the following table).

	Bayesian	Regret-based
Value of a Single Recommendation	Expected Utility	Minimax Regret
Value of a Recommendation set	Expected Utility of a Selection (EUS)	Minimax *Setwise* Regret
Value of a Query	Expected posterior Utility (EPU)	Worstcase Regret (WR)

Bayesian Elicitation The system's uncertainty about the user preferences is reflected in a distribution, $P(w; \theta)$ over the space W of possible utility functions. Here θ denotes the parameterization of our model, and we often refer to θ 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 θ, the optimal option $x^*(\theta)$ is that with greatest expected utility $EU^*(\theta) = \max_{x \in \mathbf{X}} EU(x; \theta)$, with $x^*(\theta) = \arg\max_{x \in \mathbf{X}} EU(x; \theta)$. When the user selects an option x in a choice set S, the belief is updated to $P(w; \theta | S \rightsquigarrow x)$.

In broad terms, we assume that the utility of a recommendation set S is the utility of its most preferred item. A probabilistic model can account for "noisy" responses: when presented with a choice query, the user will, with some probability, choose something different than her true preferred item. 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. Under utility uncertainty, probability of selection/response is given by $P_R(S \rightsquigarrow x_i; \theta) = \int_W P_R(S \rightsquigarrow x_i; w)P(w; \theta)dw$. We then define the *expected utility of selection (EUS)* of recommendation set S given θ and R:

$$EUS_R(S; \theta) = \sum_{x \in S} P_R(S \rightsquigarrow x; \theta) EU(x; \theta | S \rightsquigarrow x) \quad (1)$$

We consider different response models. In the *noiseless response model*, $P_R(S \rightsquigarrow x; w) = 1$ if $w \in S \triangleright x_i$, 0 otherwise. The *constant noise model* instead assumes each option x, apart from the most true preferred option, is selected with (small) constant probability, independent of w. Finally in the *logistic* response model, commonly used in choice modeling, 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 γ is a temperature parameter.[2]

When treating S as a query set (as opposed to a recommendation set), we are not interested in its expected utility, but rather in its *expected value of information (EVOI)*, or the (expected) degree to which a response will increase the quality of the system's recommendation. Given belief state θ, the *expected posterior utility* (*EPU*) of query set S under response model R is

$$EPU_R(S; \theta) = \sum_{x \in S} P_R(S \rightsquigarrow x; \theta) \, EU^*(\theta | S \rightsquigarrow x) \quad (2)$$

$EVOI(S; \theta)$ is then $EPU(S; \theta) - EU^*(\theta)$, the expected improvement in decision quality given S. An optimal query (of fixed size k) is any S with maximal $EVOI$, or equivalently, maximal EPU.

Regret-based Elicitation The system maintains an explicit representation of the set W^θ of feasible utility functions (as in the previous case, θ represents the belief state) that are consistent with previous responses to queries (and possibly prior knowledge); we refer to as *strict uncertainty*.

[2]For comparison queries (i.e., $|S| = 2$), P_L is the logistic function of the difference in utility between the two options. This model is also variously known as the Luce-Sheppard, Bradley-Terry, or *mixed multinomial logit*.

Unlike the Bayesian case, we do not have probabilistic information about the relative likelihood of the different $w \in W^\theta$. When the utility function is linear in its bounded parameters w, the belief θ is characterized by a set of linear constraints that define W^θ to be a convex polytope. Whenever the user answers a query, the polytope gets trimmed; "belief update" in this setting consists of simply adding the new constraints: $W^{\theta|S \rightsquigarrow x_i} = W \cap S \triangleright x_i$.

Minimax regret (Savage 1954; Kouvelis and Yu 1997; Boutilier et al. 2006) can be used to generate recommendations that minimize the worst-case loss incurred by assuming an adversary will choose the user's utility function w from W^θ to maximize the difference in utility between the optimal configuration (under w) and the recommendation. The maximum regret $MR(x;\theta)$ of choosing x is $MR(x;\theta) = \max_{x^a \in \mathbf{X}} \max_{w \in W^\theta} u(x^a;w) - u(x;w)$; the *minimax-regret optimal* recommendation is the item minimizing max-regret: $x^*(\theta) = \arg\min_{x \in \mathbf{X}} MR(x;\theta)$. Minimax regret $MMR(\theta)$ is its value. $MR(x;\theta)$ bounds the loss associated with x, and is zero iff x is optimal for all $w \in W^\theta$; any choice that is not minimax optimal has strictly greater loss than $x^*(\theta)$ for some $w \in W^\theta$.

We now generalize minimax regret to *recommendation sets*. Define the *setwise max regret* of a set S:

$$SMR(S;\theta) = \max_{x^a \in \mathbf{X}} \max_{w \in W^\theta} \quad u(x^a;w) - \max_{x \in S} u(x;w) =$$
$$= \max_{x^a \in \mathbf{X}} \max_{w \in W^\theta} \min_{x \in S} \; u(x^a;w) - u(x;w)$$

Intuitively, given a recommendation set, an adversary wanting to maximize regret should do so assuming that the user can select *any* from a set of k options. Formally, we choose the set of k options first, but delay the specific choice from the slate until *after* the adversary has chosen a utility function w. The regret of a set is the difference between the utility of the best configuration under w and the utility of the best option w.r.t. w in the slate.

When treating S as a query set (as opposed to a recommendation set), we are not interested in its max regret, but rather in *how much a query response will reduce minimax regret*. In our distribution-free setting, the most appropriate measure is *myopic worst case regret (WR)*, a measure of the value of information of a query. Generalizing the pairwise measure of (Boutilier et al. 2006), we define the *myopic worst-case regret (WR)* of $S = \{x_1, \ldots, x_k\}$ to be:

$$WR(S;\theta) = \max[MMR(\theta|S \rightsquigarrow x_1), \ldots, MMR(\theta|S \rightsquigarrow x_k)]$$

An *optimal query set* S^* is any S (of fixed size) that minimizes worst case regret. Intuitively, each possible response $x_i \in S$ gives rise to updated beliefs about the user's utility function. We use the worst-case response to measure the quality of the query (i.e., the response that leads to the updated W with greatest remaining minimax regret). The optimal query minimizes this value.

Theoretical Results

We now develop the connection between optimal recommendation sets and optimal choice queries. As our theoretical results are formally analogous for both uncertainty representations (Bayesian and regret-based), we present them together. With a slight abuse of notation, let θ denote the system's current knowledge about user preferences. In the regret-based model, θ is a set of constraints on utility parameters; in the Bayesian model, θ reflects probabilistic beliefs about the value of utility parameters.

First notice that, as a direct consequence of our definitions, the value of a set as recommendation *bounds* its value as a choice query. More formally, $EPU(S;\theta) \geq EUS(S;\theta)$ (for any response model). In the regret-based approach, as the objective is to be minimized, $WR(S;\theta) \leq SMR(S;\theta)$. We can then introduce a transformation T_θ that modifies a set S such that its value as recommendation set usually increases. This transformation is used in two ways: (i) to prove the optimality of optimal recommendation sets when used as query sets; (ii) and directly as a computationally viable heuristic strategy for generating query sets.

Definition 1 *Let $S = \{x_1, \cdots, x_k\}$ be a set of options. Define: $T_\theta(S) = \{x^*(\theta|S \rightsquigarrow x_1), \cdots, x^*(\theta|S \rightsquigarrow x_k)\}$where $x^*(\theta|S \rightsquigarrow x_i)$ is the optimal option when the belief θ is conditioned on $S \rightsquigarrow x_i$.*

Intuitively, T refines a recommendation set S of size k by producing k updated beliefs based on each possible user choice from S, and replacing each option in S with the optimal option under the corresponding update (where "optimal" is either in expectation or minimax-regret optimal depending on the model of uncertainty).

We consider noiseless responses first. In this case, T improves the value of set S in the following sense: *the value of the resulting set $T_\theta(S)$ as a recommendation is no less than the value of the original set S as query*. Formally:

- In the Bayesian approach, $EUS(T_\theta(S);\theta) \geq EPU(S;\theta)$ (Viappiani and Boutilier 2010)
- In the regret-based model, $SMR(T_\theta(S), \theta) \leq WR(S, \theta)$ (Viappiani and Boutilier 2009)

The proof is quite technical, but relies on partitioning the space W based on $S \triangleright x_i$ and $T(S) \triangleright x_i$. From this observation (which we dub the *query iteration lemma*) our main results follow. Consider the optimal recommendation set (i.e., maximizing EUS in the Bayesian model, or minimizing SMR in the regret-based model). While one might suppose that a weaker recommendation set could yield better value as a query, our results show that this is not possible: the function T provides an operational mechanism for producing a set that is optimal as both a recommendation set and as a choice query. Somewhat surprisingly, there is no exploration/exploitation tradeoff after all! The decision-theoretic "diversity" of the optimal recommendation set is (myopically) maximally informative for elicitation.

Theorem 1 *Let S^* be an optimal recommendation set of size k. Then S^* is an optimal choice query. More formally,*

- *In the Bayesian model, let $S^* = \arg\max_S EUS(S, \theta)$. Then $EPU(S^*, \theta) = \max_S EPU(S)$*
- *Let $S^* = \arg\min_S SMR(S, \theta)$ in the regret-based model. Then $WR(S^*_W, \theta) = \min_S WR(S;\theta)$*

Furthermore, the value of the optimal set as a recommendation coincides with its value as query: $EPU(S^*, \theta) = EUS(S^*, \theta)$ *and* $WR(S^*; \theta) = SMR(S^*; \theta)$.

The fact that the same set of options comprises the optimal recommendations and the (myopically) optimal choice query means that we can leave termination decisions in the hands of the user without relying on distinct decision/querying phases.

Noisy Response Models We have analyzed the impact of several noisy response models in the Bayesian setting (regret-based elicitation does not directly account for noise). For the constant noise model, the query iteration lemma holds and the optimal recommendation set of size at most k is also the optimal query set of size at most k.[3] For the logistic model, we analyze the loss $\Delta(S; \theta) = EUS_{NL}(S; \theta) - EUS_L(S; \theta)$ that arises due to noise. In this model, selection probabilities are function of the difference of utilities between items and there is a precise value for which loss is maximum. Based on this, we show that optimal recommendation sets are near-optimal queries. We also show that the bound is surprisingly small in practice (Viappiani and Boutilier 2010).

Algorithms

Our theoretical results also have practical implications. Optimizing a recommendation set is a simpler problem with lower complexity than typical methods for computing EVOI. So even in settings where elicitation is strictly separated from recommendation, one can optimize w.r.t. the simpler objective in both cases. Optimization of (Bayesian) recommendation sets is also *submodular*, so greedy optimization gives strong worst-case guarantees and can be implemented using lazy evaluation to achieve a significant speedup (Krause and Guestrin 2007). Moreover, noiseless optimization is quite effective even when evaluated in a noisy setting, as guaranteed by our analysis of the logistic response model.

The T transformation gives rise to a natural heuristic method for computing good query/recommendation sets. *Query iteration (QI)* starts with an initial set S, and locally optimizes S by repeatedly applying operator $T(S)$ until a fixed point is reached.[4] While the optimal set is a fixed point, there may be others; thus, QI is sensitive to the initial set S. We consider several initialization strategies. In the regret-based approach, the *current solution*, consisting of the MMR-optimal x^* and the adversarial choice x^a, can be used to initialize QI; and the approach can be generalized to sets of any size k (Viappiani and Boutilier 2009).

[3] This weaker form of the Theorem is due to the fact that T might produce a set of smaller cardinality (this is not a concern in the noiseless model, as EUS is monotone). We thank the anonymous reviewer for this observation. If one allows multisets (the same option can be shown more than once) then Theorem 1 holds.

[4] In the Bayesian model, such fixed point corresponds to the condition $EUS(T(S); \theta)$=$EUS(S; \theta)$; the condition is $SMR(T(S); \theta)$=$SMR(S; \theta)$ in the regret-based approach.

Experimental results (Viappiani and Boutilier 2009; 2010) involving decision problems with hundreds of options show that our strategies are very effective in generating queries that quickly reduce the *actual* regret or loss. In most simulations, the optimal product is discovered with just a few queries. QI is extremely fast, computing optimal sets in a fraction of a second, and is particularly well-suited to large datasets (see the following table for an experiment with Bayesian elicitation).

Computation Time	exact EPU	exact EUS	greedy EUS with lazy eval	Query Iteration
Dataset 1 Size=187	1815s	405s	1.02s	0.15s
Datase 2 Size=506	2 weeks	2 days	0.93s	0.05s

Future Directions

We are interested in elicitation strategies that combine probabilistic and regret-based models. Other interesting directions are: further theoretical and practical investigation of local search strategies such as query iteration, the development of strategies for elicitation in large-scale configuration problems, and automatically learning user response models.

References

Boutilier, C.; Patrascu, R.; Poupart, P.; and Schuurmans, D. 2006. Constraint-based optimization and utility elicitation using the minimax decision criterion. *Artifical Intelligence* 170(8–9):686–713.

Boutilier, C.; Zemel, R. S.; and Marlin, B. 2003. Active collaborative filtering. *19th Conf. on Uncertainty in AI (UAI-03)*, 98–106.

Boutilier, C. 2002. A POMDP formulation of preference elicitation problems. *18th National Conf. on AI (AAAI-02)*, 239–246.

Braziunas, D., and Boutilier, C. 2008. Elicitation of factored utilities. *AI Magazine* 29(4):79–92.

Chajewska, U.; Koller, D.; and Parr, R. 2000. Making rational decisions using adaptive utility elicitation. *17th National Conf. on AI (AAAI-00)*, 363–369.

Kouvelis, P., and Yu, G. 1997. *Robust Discrete Optimization and Its Applications*. Dordrecht: Kluwer.

Krause, A., and Guestrin, C. 2007. Near-optimal observation selection using submodular functions. *22nd National Conf. on AI (AAAI-07)*, 1650–1654.

Louviere, J., Hensher, D., and Swait, J. 2000. *Stated Choice Methods: Analysis and Application*. Cambridge Univ. Press.

Price, R., and Messinger, P. R. 2005. Optimal recommendation sets: Covering uncertainty over user preferences. *20th National Conf. on AI (AAAI-05)*, 541–548.

Reilly, J.; McCarthy, K.; McGinty, L.; and Smyth, B. 2005. Incremental critiquing. *Knowledge-Based Systems* 18(4–5):143–151.

Reilly, J.; Zhang, J.; McGinty, L.; Pu, P.; and Smyth, B. 2007. Evaluating compound critiquing recommenders: a real-user study. *8th ACM Conf. on Electronic Commerce (EC-07)*, 114–123.

Savage, L. 1954. *The Foundations of Statistics*. NYC: Wiley.

Viappiani, P., and Boutilier, C. 2009. Regret-based optimal recommendation sets in conversational recommender systems. *3rd ACM Conference on Recommender Systems (RecSys09)*, 101–108.

Viappiani, P., and Boutilier, C. 2010. Optimal Bayesian recommendation sets and myopically optimal choice query sets. In *Advances in Neural Information Processing Systems 23 (NIPS)*, 2352–2360.