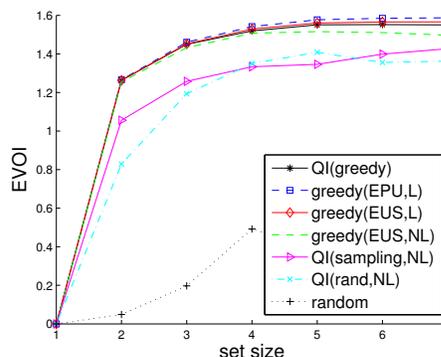
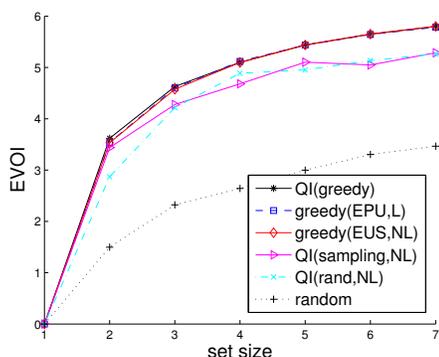

Optimal Bayesian Recommendation Sets and Myopically Optimal Choice Query Sets: Supplementary Material

Paolo Viappiani*
Department of Computer Science
University of Toronto
Toronto, ON, CANADA
paolo@cs.toronto.edu

Craig Boutilier
Department of Computer Science
University of Toronto
Toronto, ON, CANADA
cebly@cs.toronto.edu

Appendix A: Additional Experiments



(a) EVOI wrt set size (187 outcomes, 30 runs, R_{NL}) (b) EVOI wrt set size (506 outcomes, 30 runs, R_L)

We analyze the impact of the size of the query set, varying k from 2 to 7. In Figure 1(a) we plot expected value of information of a query as a function of the number of items in the first dataset, assuming a noiseless response model, for few of the strategies considered before. Greedy maximization of EPU (more computationally demanding) shows no advantage over EUS maximization. Similarly in Figure 1(b) we consider the second dataset, this time in a scenario of increased noise with $\gamma = 0.33$; even in this case the difference between maximization of EPU and EUS is minimal, and discernible only for $k \geq 4$.

Appendix B: Proofs

In this appendix we provide all proofs for our theoretical results. First, we clarify part of the notation. We recall that $S \triangleright x$ is the event “ x is the item with highest utility in the set S ” and regions $W \cap S \triangleright x_i$, $x_i \in S$ partition utility space. We call this the S -partition of W . We consider $\theta[S \triangleright x]$ to be the normalized projection of the belief θ on $W \cap S \triangleright x$, whose value is $\frac{1}{Z}P(w)$ when $w \in W \cap S \triangleright x$, and 0 otherwise (where Z is a normalizing constant). We also refer with $P(S \triangleright x; \theta)$ to the probability mass of such projection under belief θ (the probability, according to belief θ , that x is the true best option in the slate).

*From 9/2010 to 12/2010 at the University of Regina; from 01/2011 onwards at Aalborg University.

$$p(w; \theta[S \triangleright x]) = \begin{cases} \frac{P(w; \theta)}{P(S \triangleright x; \theta)} & \text{if } w \in W \cap S \triangleright x \\ 0 & \text{otherwise} \end{cases} \quad \text{where } P(S \triangleright x; \theta) = \int_{W \cap S \triangleright x} P(w; \theta) dw \quad (1)$$

Proof of Equation 3:

$$EUS_R(S; \theta) = \sum_{x \in S} P(S \rightsquigarrow x; \theta) \int_W u(x; w) P(w; \theta | S \rightsquigarrow x) dw = \sum_{x \in S} P(S \rightsquigarrow x; \theta) \int_W u(x; w) \frac{P(S \rightsquigarrow x | w; \theta)}{P(S \rightsquigarrow x; \theta)} P(w) dw \quad (2)$$

$$= \sum_{x \in S} \int_W u(x; w) P_R(S \rightsquigarrow x; w) P(w; \theta) dw = \int_W \sum_{x \in S} [P_R(S \rightsquigarrow x; w) u(x; w)] P(w; \theta) dw \quad (3)$$

where we expanded the expression of the expected utility in the conditioned space, used Bayes theorem (2), substituted the response model and moved the summation inside the integration sign. Note that $P(S \rightsquigarrow x | w; \theta) = P_R(S \rightsquigarrow x; w)$ as the response model is exactly defined as the likelihood of a user selection response given the utility weights. ■

Proof of Theorem 1: We consider the gain in *EUS* that sets S and Q , with $S \subseteq Q$, incur by the addition of a given item $\{a\}$. Let's call $Gain(S) = EUS(\{a\} \cup S) - EUS(S)$ (for ease of notation, we make implicit the current belief θ in the expressions of this proof) and consider $Gain(S, w)$, the gain contribution for each w (cfr. Observation 3) so that $Gain(S) = \int_w p(w) Gain(S; w) dw$. It follows:

$$Gain(S; w) = P_R(S' \rightsquigarrow a; w) u(a; w) + \sum_{x \in S} [P_R(S' \rightsquigarrow x; w) - P_R(S \rightsquigarrow x; w)] u(x; w) \quad (4)$$

where $S' = S \cup \{a\}$ and $Q' = Q \cup \{a\}$. We will show that $Gain(S) \geq Gain(Q)$ (characteristics of diminishing returns) by verifying that, for each w , we have $Gain(S; w) \geq Gain(Q; w)$. We do that separately for each response model (NL, C).

Noiseless response model (NL):

Recall that $x_w^*(S)$ indicates the (true) optimal product in set S given utility weights w . We consider different cases with respect to the optimality of item $\{a\}$ in the set $S' = S \cup \{a\}$ and $Q' = Q \cup \{a\}$.

First, if w is such that a optimal in S' and optimal in Q' , $Gain(S', w) = u(a; w) - u_w^*(S) \geq u(a; w) - u_w^*(Q) = Gain(Q', w)$, since $u_w^*(Q) \geq u_w^*(S)$. Second, suppose w such that a is optimal in S' , but not optimal in Q' ; then, $Gain(S', w) = u(a; w) - u_w^*(S) \geq 0 = Gain(Q', w)$. Third, suppose w such that a is neither optimal in S' , nor in Q' : the gain is zero for both sets. Finally, the following case is impossible: w such that a optimal in Q' , but not optimal in S' (impossible because $S' \subseteq Q'$).

Thus *EUS* with noiseless responses is submodular.

Constant error response model (C):

We substitute $P_R(S \rightsquigarrow x; w)$ in the expression of the gain according to the constant response model. Most terms in the summation in Equation 4 evaluate to β and cancel out. Note that the likelihood α of a correct answer is a function of the set size k (when there are more items, it is easier for the user to make a mistake). Let $\alpha_{Q'} = \alpha(|Q'|)$ (similarly for Q, S and S'), it holds $\alpha_{S'} = \alpha_S - \beta$.

Again, we consider different cases with respect to the optimality of item $\{a\}$ in the set $S' = S \cup \{a\}$ and $Q' = Q \cup \{a\}$, and write the expression of the gain in each case.

1. Suppose w is such that a optimal in S' and optimal in Q' :

$$Gain(S', w) = \alpha_{S'} [u(a; w) - u_w^*(S)] \geq 0 \quad (5)$$

$$Gain(Q', w) = \alpha_{Q'} [u(a; w) - u_w^*(Q)] \geq 0 \quad (6)$$

2. Suppose w such that a is optimal in S' , but not optimal in Q' . S has a gain; while Q a loss.

$$Gain(S', w) = \alpha_{S'} [u(a; w) - u_w^*(S)] \geq 0 \quad (7)$$

$$Gain(Q', w) = \beta [u(a; w) - u_w^*(Q)] \leq 0 \quad (8)$$

3. Suppose w such that a is neither optimal in S' , nor in Q' . The gain will be negative for both sets (thus a loss).

$$Gain(S', w) = \beta [u(a; w) - u_w^*(S)] \leq 0 \quad (9)$$

$$Gain(Q', w) = \beta [u(a; w) - u_w^*(Q)] \leq 0 \quad (10)$$

4. The following case is impossible: w such that a optimal in Q' , but not optimal in S' (impossible because $S' \subseteq Q'$).

In all the cases $\text{Gain}(S; w) \geq \text{Gain}(Q; w)$, the gain is greater for S than for the larger set Q (in the third case, S suffers a smaller loss): by componentwise comparison, and by virtue that $u_w^*(Q) \geq u_w^*(S)$ and $\alpha_{S'} > \alpha_{Q'}$ and that, for the model to be meaningful we have $\alpha > \beta$. ■

Observation 1

$$EUS_C(S, \theta) = \sum_{x \in S} P(S \triangleright x; \theta) \left\{ \alpha EU(x; \theta | S \triangleright x) + \beta \sum_{y \neq x} EU(y; \theta | S \triangleright x) \right\} \quad (11)$$

Proof of Observation 1: We use equation 3 to write the value of a set, then we decompose the integration over each element of the S -partition. Finally, we observe that in the constant error model the likelihood of the response is independent of w , but only depends on which element of the partition it belongs to.

$$EUS(S; \theta) = \int_W \sum_{y \in S} [u(y; w) P_R(S \rightsquigarrow y; w)] P(w; \theta) dw = \sum_{x \in S} \int_{W \cap S \triangleright x} \sum_{y \in S} \{u(y; w) P_R(S \rightsquigarrow y; w)\} P(w; \theta) dw = \quad (12)$$

$$= \sum_{x \in S} \left\{ \alpha \int_{W \cap S \triangleright x} u(x; w) P(w; \theta) dw + \beta \sum_{y \in S, y \neq x} \int_{W \cap S \triangleright x} u(y; w) P(w; \theta | S \triangleright x) dw \right\} \quad (13)$$

$$= \sum_{x \in S} P(S \triangleright x; \theta) \left\{ \alpha EU(x, \theta | S \triangleright x) + \beta \sum_{y \neq x} EU(y, \theta | S \triangleright x) \right\} \quad (14)$$

■

Proof of Lemma 1: Let $S = \{x_1, \dots, x_k\}$ be a set of options, and $T_R(S) = \{x'_1, \dots, x'_k\}$ be the set resulting from the application of the transformation T . For R_{NL} , the **noiseless** response model, the argument relies on partitioning W w.r.t. options in S :

$$EPU_{NL}(S; \theta) = \sum_{i,j} P(S \triangleright x_i, T(S) \triangleright x'_j; \theta) EU(x'_i, \theta | S \triangleright x_i, T(S) \triangleright x'_j) \quad (15)$$

$$EUS_{NL}(T(S); \theta) = \sum_{i,j} P(S \triangleright x_i, T(S) \triangleright x'_j; \theta) EU(x'_j, \theta | S \triangleright x_i, T(S) \triangleright x'_j) \quad (16)$$

Compare the two expression componentwise: 1) If $i = j$ then the components of each expression are the same. 2) If $i \neq j$, for any w with nonzero density in $\theta[S \triangleright x_i, T(S) \triangleright x'_j]$, we have $u(x'_j; w) \geq u(x'_i, w)$, thus $EU(x'_j) \geq EU(x'_i)$ in the region $S \triangleright x_i, T(S) \triangleright x'_j$. Since $EUS_{NL}(T(S); \cdot) \geq EPU_{NL}(S; \cdot)$ in each component, the result follows.

For the **constant error response model** C , we use the observations that EPU and EUS can be expressed in function of the S -partition. Call $\lambda_{i,j} = P(S \triangleright x_i, T(S) \triangleright x'_j; \theta)$ the probability of being in the space where x_i is the best item in slate S and x'_j is the best in slate $T(S)$, given the current belief θ .

$$EPU(S; \theta) = \sum_{i,j} \lambda_{i,j} \left\{ \alpha EU(x'_i, \theta | S \triangleright x_i, T(S) \triangleright x'_j) + \beta \sum_{o \neq i} EU(x'_o, \theta | S \triangleright x_i, T(S) \triangleright x'_j) \right\} \quad (17)$$

$$EUS(T(S); \theta) = \sum_{i,j} \lambda_{i,j} \left\{ \alpha EU(x'_j, \theta | S \triangleright x_i, T(S) \triangleright x'_j) + \beta \sum_{o \neq j} EU(x'_o, \theta | S \triangleright x_i, T(S) \triangleright x'_j) \right\} \quad (18)$$

As before, we compare $EPU(S; \theta)$ and $EUS(T(S); \theta)$ componentwise to show that the latter is greater:

- If $i = j$ then the expressions within the brackets give the same results
- If $i \neq j$ then $EU(x'_j, \theta | S \triangleright x_i, T(S) \triangleright x'_j) \geq EU(x'_i, \theta | S \triangleright x_i, T(S) \triangleright x'_j)$ by virtue of the projection, it holds $T(S) \triangleright x'_j$, so x'_j has higher utility than x'_i by definition. Note also that the two expressions are convex combinations of the expected utilities of the same items in $T(S)$ wrt the projected beliefs in the $T(S)$ -partition. It follows that, if $\alpha \geq \beta$ the component of $EUS(T(S); \theta)$ is greater (or equal) than the component of $EPU(S; \theta)$.

■

Proof of Theorem 2: Suppose S^* is not an optimal query set, i.e., there is some S s.t. $EPU(S; \theta) > EPU(S^*; \theta)$. Applying T to S gives a new query set $T(S)$, which by the results above shows: $EUS(T(S); \theta) \geq EPU(S; \theta) > EPU(S^*; \theta) \geq EUS(S^*; \theta)$. This contradicts the EUS-optimality of S^* .

■

Proof of Theorem 3:

We first consider the case $k = 2$ (pairs of items). As discussed in the paper, the value of the maximal loss Δ_{\max} is function only of the difference in utility of two options. For a specific value of $z \geq 0$, EUS-loss is exactly the utility difference z times the probability of choosing the less preferred option under R_L : $1 - L(\gamma z) = L(-\gamma z)$ where L is the logistic function.

$$\Delta_{\max} = \max f(z) = \max \frac{z}{1 + e^{\gamma z}} \quad (19)$$

We impose the derivative equal to zero:

$$\frac{\partial f}{\partial z} = 0 \Leftrightarrow \frac{1}{1 + e^{\gamma z}} + z \frac{-e^{\gamma z}}{(1 + e^{\gamma z})^2} = 0 \Leftrightarrow \frac{1}{1 + e^{\gamma z}} \left(1 - \gamma z \frac{e^{\gamma z}}{1 + e^{\gamma z}} \right) = 0 \Leftrightarrow 1 + e^{\gamma z} - \gamma z e^{\gamma z} = 0 \Leftrightarrow \quad (20)$$

We solve the equation in z :

$$(\gamma z - 1)e^{\gamma z} = 1 \Leftrightarrow (\gamma z - 1)e^{\gamma z - 1} = e^{-1} \Leftrightarrow \gamma z - 1 = \mathcal{L}_W\left(\frac{1}{e}\right) \quad (21)$$

where $\mathcal{L}_W(\cdot)$ is the Lambert-W function. Moreover, the last expression of Eq. 20, substituted into Eq. 19, gives $\Delta_{\max} = f(z_{\max}) = \frac{e^{-\gamma z_{\max}}}{\gamma} = z_{\max} - \frac{1}{\gamma}$. Thus:

$$z_{\max} = \frac{1}{\gamma} \left[1 + \mathcal{L}_W\left(\frac{1}{e}\right) \right] \quad ; \quad \Delta_{\max} = \frac{1}{\gamma} \mathcal{L}_W\left(\frac{1}{e}\right) \quad (22)$$

The argument is similar for $k = 3$. Given three options, x_1, x_2 and x_3 , we define $z_{i,j} = u(x_i) - u(x_j)$ to be the difference in utility between two options. Assuming, without loss of generality, that x_1 is the utility maximizing option in the set ($S \triangleright x_1$), the loss function is the following:

$$f(z_{1,2}, z_{1,3}, z_{2,3}) = z_{1,2} \frac{1}{1 + e^{z_{1,2}} + e^{-z_{2,3}}} + z_{1,3} \frac{1}{1 + e^{z_{1,2}} + e^{z_{2,3}}} \quad (23)$$

We maximize the loss by imposing $\frac{\partial f}{\partial z} = 0$; it is possible to show that $z_{1,2} = z_{1,3}$ and $z_{2,3} = 0$. The expression becomes an equation in a single variable; we let $z = z_{1,2}$ and have to solve $\gamma z - 1 - 2e^{-\gamma z} = 0$, giving $z_{\max} = \left[1 + \mathcal{L}_W\left(\frac{2}{e}\right) \right]$.

For sets of any size, once again the loss is maximized when all items beside the most preferred have the same utility; call z the difference in utility. The function to maximize is: $f(z) = z(k-1) \frac{1}{(k-1) + e^z}$ from which follows

$$z_{\max} = \frac{1}{\gamma} \left[1 + \mathcal{L}_W\left(\frac{k-1}{e}\right) \right] \quad ; \quad \Delta_{\max} = \frac{1}{\gamma} \mathcal{L}_W\left(\frac{k-1}{e}\right) \quad (24)$$

■

Proof of Lemma 2: Let $x'_i = x^*(\theta|S \rightsquigarrow x_i)$ under R_L . The expressions of $EUS_{NL}(T(S))$ and $EPU_L(S)$ can be rearranged in the following way:

$$EUS_{NL}(T_L(S); \theta) = \sum_i \int_{W \cap T(S) \triangleright x'_i} u(x'_i; w) P(w; \theta) dw \quad (25)$$

$$EPU_L(S; \theta) = \sum_i \int_{W \cap T(S) \triangleright x'_i} \sum_j P_R(S \rightsquigarrow x_j; w) u(x'_j; w) P(w; \theta) dw \quad (26)$$

$$(27)$$

We compare the two expressions componentwise. In the partition $W \cap S \triangleright x_i$, x'_i is the best item in the slate $T(S)$, giving higher utility than any other x'_i with $j \neq i$. Therefore $u(x'_i; w)$ is greater than any convex combination of the (lower or equal) values $u(x'_j; w)$. Thus $EUS_{NL}(T_L(S))$ is greater. ■

Proof of $EPU_{NL}^*(\theta) \geq EPU_L^*(\theta)$ (consequence of Lemma 2): Let q_L^* be the optimal query set with respect the current belief θ and the logistic response model: $q_L^* = \arg \max_q EPU_L(q; \theta)$ and $EPU_L^*(\theta) = EPU_L(q_L^*; \theta)$, we derive (we drop parametrization wrt θ in the following):

$$EPU_{NL}^* = EUS_{NL}^* \geq EUS_{NL}(T_L(q_L^*)) \geq EPU_L(q_L^*) = EPU_L^* \quad (28)$$

■

Proof of Theorem 4: Consider the optimal query S_L^* and the set $S' = T_L(S_L^*)$ obtained by applying T_L . From Lemma 2, $EUS_{NL}(S'; \theta) \geq EPU_L(S_L^*; \theta) = EPU_L^*(\theta)$. From Thm. 3, $EUS_L(S'; \theta) \geq EUS_{NL}(S'; \theta) - \Delta_{\max}$; and from Thm. 2, $EUS_{NL}^*(\theta) = EPU_{NL}^*(\theta)$. Thus $EUS_L^*(\theta) \geq EUS_L(S'; \theta) \geq EUS_{NL}(S'; \theta) - \Delta_{\max} \geq EPU_L^*(\theta) - \Delta_{\max}$ ■