
Characterizing EVOI-Sufficient k -Response Query Sets in Decision Problems

Robert Cohn

Computer Science and Engineering
University of Michigan
rwcohn@umich.edu

Satinder Singh

Computer Science and Engineering
University of Michigan
baveja@umich.edu

Edmund Durfee

Computer Science and Engineering
University of Michigan
durfee@umich.edu

Abstract

In finite decision problems where an agent can query its human user to obtain information about its environment *before* acting, a query’s usefulness is in terms of its Expected Value of Information (EVOI). The usefulness of a query set is similarly measured in terms of the EVOI of the queries it contains. When the only constraint on what queries can be asked is that they have exactly k possible responses (with $k \geq 2$), we show that the set of k -response decision queries (which ask the user to select his/her preferred decision given a choice of k decisions) is *EVOI-Sufficient*, meaning that no single k -response query can have higher EVOI than the best single k -response decision query for any decision problem. When multiple queries can be asked before acting, we provide a negative result that shows the set of depth- n query trees constructed from k -response decision queries is not EVOI-Sufficient. However, we also provide a positive result that the set of depth- n query trees constructed from k -response *decision-set* queries, which ask the user to select from among k sets of decisions as to which set contains the best decision, is EVOI-Sufficient. We conclude with a discussion and analysis of algorithms that draws on a connection to other recent work on decision-theoretic knowledge elicitation.

1 Introduction

An agent acting autonomously on behalf of its user in a complex environment might have uncertainty in its environment model for a variety of reasons. For example, its

model might be incomplete (the user omitted information about circumstances not expected to arise), imprecise (user abstracted away details to simplify the modeling process), and/or outdated (user has not maintained the model to reflect current conditions). We consider a Bayesian setting where such an agent is faced with making a single decision (such as which policy to follow in the future) under uncertainty over which of a space of possible models correctly represents its environment, i.e., each candidate model prescribes utility/value to each possible decision. The best decision the agent can make, then, is one that achieves the highest value in expectation over its uncertainty.

Such an agent can benefit by *querying* its user to learn more about its environment and user’s preferences before making its decision. For instance, a robotic car choosing a route might have uncertainty about its user’s current valuation of speed versus money as it contemplates taking a toll road. By actively querying its user for information (e.g., by asking a query like “For this trip, is it better to arrive 10 minutes earlier or to save \$1?”), the agent can develop a more complete, precise, and updated preference-model and thereby behave to best suit its user’s preferences.

In this paper, we study the question of what the agent should ask its user *before* making its decision when (1) the agent may ask only a single query (myopic query selection); and (2) the agent may ask n queries (nonmyopic query selection). In our Bayesian decision making setting, the agent’s goal is to act so as to maximize its expected value; thus, the agent should ask queries that maximally improve its ability to do so. The natural criterion for evaluating a query, then, is its Expected Value of Information (EVOI), which measures the expected increase in value associated with adapting the agent’s decision as a function of the user’s response to the query.

The problem of selecting one or more queries from some query set so as to maximize EVOI (myopically or nonmyopically) has been widely studied. However, there has been little work on how to design a good query set in the first place. In this paper, we examine how to design a good query set that could contain any queries among the set of all k -response queries (queries with k possible responses).

Appearing in Proceedings of the 17th International Conference on Artificial Intelligence and Statistics (AISTATS) 2014, Reykjavik, Iceland. JMLR: W&CP volume 33. Copyright 2014 by the authors.

We show that in myopic settings, where the goal is to select a query so as to maximize EVOI without considering any future queries that could be asked, the set of k -response decision queries (a decision query asks for the best decision out of some subset) is sufficiently general in that there is no benefit in considering any k -response queries beyond k -response decision queries. This result dovetails with recent work by Viappiani and Boutilier (2010), who contribute efficient algorithms for k -response decision query selection enjoying approximation guarantees due to the submodularity of myopically optimal decision query EVOI.

In addition, we show that in a nonmyopic setting, where the goal is to select a depth- n k -response query tree instead of a single query, the set of depth- n trees constructed from k -response decision queries is not sufficiently general, while the set of depth- n trees constructed from k -response decision-*set* queries is. Finally, we show the computational result that depth- n k -response query tree selection can be reduced to k^n -response decision query selection, where the algorithms contributed by Viappiani and Boutilier (2010) directly apply.

2 Related Work

A rich literature exists on the subject of querying and information acquisition in non-decision-theoretic settings and/or under different query selection criteria than ours, appearing in such areas as Bayesian Experimental Design (Chaloner and Verdinelli, 1995), Active Learning (Settles, 2009; Nowak, 2011), Preference Elicitation (Braziunas and Boutilier, 2008; Regan and Boutilier, 2009), Human-robot/agent Interaction (Cakmak et al., 2010), and Optimal Learning (Powell and Frazier, 2008). We here however restrict our comparison to related work studying decision-theoretic settings that use EVOI to measure query informativeness. Within this space, our approach can be compared to related work along the dimensions that follow.

Source of query set.

In scoping what queries the agent can consider asking, researchers typically provide the agent with a specification of the set of possible queries, fitted to the agent’s setting. In contrast, our goal is to provide insight on what constitute good query sets. Thus, we consider more general query sets that are only constrained by the number of possible responses that queries can have, rather than being predisposed to querying about particular aspects such as parameter values or decisions.

Generality in structure of agent uncertainty and decision problem.

Various authors have studied queries under specific structures of uncertainty and/or decision problems, e.g., uncertainty in parameters defining goals/rewards (Chajewska et al., 2000) or system dynamics (Cohn et al., 2010). Com-

binning such structure with assumptions regarding the family of distributions expressing the agent’s uncertainty can afford fruitful computational advantages and/or theoretical results (Braziunas and Boutilier, 2005). Since our setting is general in all of these aspects, our theoretical results and algorithms cannot exploit such structure. The power of this generality, though, is that our results apply to all such settings.

Myopic and nonmyopic query selection.

Settings in which multiple queries may be asked induce a challenging optimization problem in which EVOI-based query selection involves taking into account how useful a query would be when combined with possible future queries (Boutilier, 2002). This makes for difficult sequential decision making problems because of a combination of (1) large or infinite query sets (analogous to actions) and their induced posterior distribution spaces (analogous to a state space); and (2) computationally demanding Bayesian inference for individual query response updates which itself typically requires approximations (Doshi et al., 2012).

Thus, a common approximation to this problem is to repeatedly select which single query to ask, without taking into account future queries (Dittmer and Jensen, 1997; Bayer-Zubek, 2004; Cohn et al., 2011), which is termed *myopic* query selection, in contrast to *nonmyopic* query selection. In some settings, myopic query selection offers powerful approximation guarantees (Golovin and Krause, 2011), but in fact is itself nontrivial in many contexts (Braziunas and Boutilier, 2008), especially when evaluating even a single query is expensive (Cohn et al., 2011; Wilson et al., 2012). We consider both myopic and nonmyopic query selection in Sections 5 and 6, respectively.

Querying Process.

We assume the agent is acting in a setting where a decision is to be made after querying ends (e.g., adopting a policy (Chajewska et al., 2000) or selecting an investment (Athey and Levin, 2001)). This is simpler than settings where queries are interleaved with other types of actions, a subject of recent interest (Doshi et al., 2012; Wilson et al., 2012); investigating extensions of our results to this setting is future work.

3 Problem Formulation

Here we define the general decision making under uncertainty problem faced by the agent, the expected value of information criterion for selecting queries, and the semantics of k -response queries that are at the core of this paper.

3.1 Decision Making under Uncertainty

We assume that the agent’s uncertainty takes the form of some arbitrary distribution ψ over a (possibly continuous)

parameter space Ω , where each $\omega \in \Omega$ specifies a possible model. Examples include uncertainty over parameter values in parameterized reward functions or transition functions, as well as other noise parameters in decision problems. We will assume that the agent is faced with selecting from a finite set U of possible decisions, where each decision could for instance be an action, an open-loop plan, or a policy. For each $\omega \in \Omega$ and decision $u \in U$, there is an associated value denoted V_ω^u that captures the utility of decision u in model ω . Then, the expected value of decision u under distribution ψ can be written as

$$V_\psi^u = \mathbb{E}_{\omega \sim \psi}[V_\omega^u] = \int_{\Omega} \psi(\omega) V_\omega^u d\omega,$$

where $\omega \sim \psi$ denotes model-parameters (ω) sampled from distribution ψ over Ω . The Bayes-optimal decision that maximizes expected value under ψ is

$$u_\psi^* = \arg \max_{u \in U} \{V_\psi^u\},$$

and the associated Bayes-optimal value is

$$V_\psi^* = \max_{u \in U} \{V_\psi^u\} = V_\psi^{u_\psi^*}.$$

3.2 Querying

Note that we have purposely not committed to how an agent computes the value V_ω^u of decision u in some model ω , because our focus here is on characterizing query sets purely in terms of their EVOI and not in terms of special structure in the models and decisions. Consequently, ψ is a sufficient statistic of the agent’s knowledge about which decision is best. Recall that in our setting the agent can pose queries to improve this knowledge. Thus, any information relevant to the value of decisions that the agent receives as a response to a query can be incorporated via a Bayes update to ψ . Specifically, assume the agent poses query q to the user, and the user responds with the j^{th} out of a finite number of possible responses to q . Then the posterior distribution of ψ given response j to query q is denoted $\psi|q=j$ with

$$\psi(\omega|q=j) = \frac{Pr(q=j|\omega)\psi(\omega)}{\int_{\Omega} Pr(q=j|\omega')\psi(\omega')d\omega'}. \quad (1)$$

Note that the likelihood function associated with a query q , $Pr(q=j|\omega)$ for each response j , completely defines the semantics of q . Intuitively, this is the agent’s model for how the user would respond to the query conditioned on parameter ω . Throughout the paper, we will use $\psi|Y=y$ to denote the posterior distribution induced when ψ is updated to incorporate the knowledge that $Y=y$ for a random variable Y . Then, after receiving response j for query q , the agent’s best decision based on the updated knowledge is $u_{\psi|q=j}^*$ with a new expected value of $V_{\psi|q=j}^*$. The expected value of asking query q , denoted $V_{\psi|q}^*$, is thus

$\mathbb{E}_{j;q,\psi}[V_{\psi|q=j}^*]$, where the expectation is over the probability of response j to query q given prior knowledge ψ . In expectation, if a query and its response induces a new Bayes-optimal decision, that decision can only improve the agent’s expected value with respect to the updated knowledge about the user’s true model:

$$\begin{aligned} V_\psi^* &= \max_u \mathbb{E}_{\omega \sim \psi}[V_\omega^u] = \max_u \left\{ \mathbb{E}_{j;q,\psi} \mathbb{E}_{\omega \sim \psi|q=j}[V_\omega^u] \right\} \\ &\leq \mathbb{E}_{j;q,\psi} \left[\max_u \mathbb{E}_{\omega \sim \psi|q=j}[V_\omega^u] \right] \text{ (by Jensen’s inequality)} \\ &= \mathbb{E}_{j;q,\psi}[V_{\psi|q=j}^*] = V_{\psi|q}^*. \end{aligned}$$

The Expected Value of Information (EVOI) associated with query q measures the expected increase in value associated with asking q and never asking another query:

$$\text{EVOI}(q, \psi) = V_{\psi|q}^* - V_\psi^*.$$

Hence, the *EVOI-optimal* query (sometimes referred to as *myopically optimal query*) from some given query set Q is given by

$$q^* = \arg \max_{q \in Q} [V_{\psi|q}^* - V_\psi^*] = \arg \max_{q \in Q} V_{\psi|q}^*.$$

We will also consider situations in which the agent can ask $n > 1$ queries. In general the choice of the next query can depend on responses to the previous queries, and so the general form to consider is a depth- n query tree. Note that greedy selection of queries based on EVOI at each node of the tree would not in general produce an EVOI-optimal depth- n query tree. This is because the effect of the resulting states of knowledge on future queries must be taken into account (we address nonmyopic query selection in Section 6).

Much past research on the subjects of selecting optimal queries and/or query policies has treated the set of queries to select from as given. For example, the query set might be chosen by the designer to simplify knowledge elicitation (e.g., queries that users find easy to answer), or to simplify knowledge updating (e.g., queries whose responses cleanly map to unambiguous updates to specific uncertain parameters). An example of each of these follows.

Action Queries are popular in Markov Decision Process (MDP) settings (e.g., in Learning by Demonstration (Chernova and Veloso, 2009) and Active Imitation Learning (Judah et al., 2012)). They typically take the form “What is the best action to do in state s ?”, and so they are designed to be easily understood and answered by users. Note that when the number of actions is finite, a query can have only a finite number of responses. Furthermore, if the number of states to query about is also finite, then the number of possible action queries is also finite.

Bound Queries ask whether the true value of some unknown parameter exceeds a particular threshold, and thus

cleanly map to the parameter uncertainty representation (e.g., (Chajewska et al., 2000; Braziunas and Bouillier, 2005)). For example, given uncertainty about the reward value of a goal state, a query might ask “Is the reward of this state above 0.5?” Bound-queries are binary-response queries in that they have two possible responses (intuitively, “yes” and “no”); however, there are an infinite number of such queries when the parameter being queried about is continuous.

3.3 k -Response Query Forms

In contrast to the work just described that emphasizes designing sets of queries that are easy to answer or easy to incorporate the answers to, we ask the question of how to design a query set that assuredly includes a query that maximizes EVOI.

First, let us consider the most valuable possible query given no constraints, i.e., when anything can be asked. Intuitively, no query can be better than one allowing the agent to behave optimally thereafter in response (i.e., allowing the agent to adopt $\max_u V_\omega^u$ under any $\omega \in \Omega$ as a function of the query response). In fact, the query that asks “What is the optimal decision?” has exactly this property, and so it has the highest possible EVOI.

Arguably, much of the burden imposed on the user in answering the above query lies in the number of responses that are possible: $|U|$. This raises the question of what form an EVOI-optimal query would take when the number of possible responses is restricted to a fixed k , corresponding to a restriction on the number of choices the user should be asked to consider. We define three types of k -response queries next.

k -Response Queries. Let Q_k denote the set of all k -response queries. Other than being limited to a finite number of responses, this class is unconstrained. For example Q_2 includes the Bound Queries mentioned previously.

k -Response Decision Queries. Let D_k denote the space of all k -response *decision queries*, where a query $q \in D_k$ asks which out of k decisions is best. As noted above, the semantics of a query are defined by its likelihood function $\Pr(q = j|\omega)$; thus the decision query q over decisions $\{u_i\}_{i=1}^k$ is defined so that $\Pr(q = j|\omega) = \delta(\arg \max_i \{V_\omega^{u_i}\} = j)$, where δ is the delta-function that takes on the value of one if the equality in its argument is satisfied and zero otherwise (we will assume that in the event of a tie, the response with smallest index is chosen).

k -Response Decision-Set Queries. Let H_k denote the set of all k -response *decision-set queries*, where a query $q \in H_k$ asks which out of k finite *sets* of decisions contains the best decision. That is, if q queries sets $\{U_i\}_{i=1}^k$, where each $U_i \subseteq U$, then $\Pr(q = j|\omega) = \delta(\arg \max_i \{\max_{u \in U_i} V_\omega^u\} = j)$ (with ties broken in the

same manner as above).

Note that the size of Q_k is infinite (in general) while the sizes of H_k and D_k are always finite in our setting. Furthermore, $Q_k \supset H_k \supset D_k$.

4 Summary of Theoretical Results

The main contributions of this paper stem from studying the following comparisons.

- Myopic query selection: Compare k -response queries with k -response decision queries in terms of the most valuable queries they contain.
- Nonmyopic query selection: Compare depth- n k -response query trees, depth- n k -response decision query trees, and depth- n k -response decision-set query trees with each other in terms of the most valuable query trees they contain.

As contributions, our comparisons show that, for all finite decision problems and for all uncertainty ψ over them,

1. (*In a myopic setting, we can safely consider only decision queries.*) We show in Section 5 that the set of k -response decision queries is *EVOI-Sufficient*: the EVOI-optimal k -response decision query has EVOI at least as high as any other k -response query.
2. (*In a nonmyopic setting, we can safely consider only decision-set queries.*) We show in Section 6 that the set of depth- n k -response decision-set query trees is *EVOI-Sufficient*: the EVOI-optimal depth- n k -response decision-set query tree has EVOI at least as high as any depth- n k -response query tree.
3. (*In a nonmyopic setting, we cannot be limited to only decision queries.*) We show in Section 6 that the set of depth- n k -response decision query trees is not *EVOI-Sufficient*: the EVOI-optimal depth- n k -response decision query tree may have lower EVOI than the EVOI-optimal depth- n k -response query tree.

5 Myopic k -Response Query Selection

When constrained to k -response queries with $k \ll |U|$, the agent will, in general, no longer have the ability to completely resolve its uncertainty regarding which decision to follow via a single query. A desirable property of a k -response query set Q , then, is that for all finite decision problems and for all uncertainty ψ over them, Q always contains a k -response query with EVOI as high as any other k -response query – i.e., when there would be no benefit in considering all k -response queries in addition to those contained by Q . We will say that such a query set is *k -response*

EVOI-sufficient (hereafter, simply *EVOI-sufficient* with the constraint k on the number of responses left implicit).

Next we present our first result, which shows that the set of k -response decision queries is EVOI-sufficient. Intuitively, the proof develops a constructive “query improvement” procedure, in that it replaces an arbitrary k -response query with a k -response decision query that has EVOI at least as high as the original query.

Theorem 1. *(The set of k -response decision queries is EVOI-Sufficient.)*

For all finite decision problems and for all uncertainty ψ over them, the EVOI-optimal k -response decision query has EVOI equal to that of the EVOI-optimal k -response query:

$$\sup_{q \in Q_k} \{EVOI(q, \psi)\} = \max_{q' \in D_k} \{EVOI(q', \psi)\}.$$

Proof. Consider an arbitrary k -response query q and recall that the decision $u_{\psi|q=j}^*$ is the Bayes-optimal decision for the posterior distribution $\psi|q = j$ (the posterior induced by the j^{th} response to query q). We will show that the k -response decision query q' that asks for the optimal decision in the set $\{u_{\psi|q=j}^*\}_{j=1}^k$ has EVOI at least as high as the EVOI of query q .

Let $Z = \{\zeta_i\}$ denote a partition of Ω such that each ζ_i has the following property: for all pairs of decisions u, u' and parameters $\omega, \omega' \in \zeta_i$, $V_{\omega'}^u(s) \geq V_{\omega'}^{u'}(s) \implies V_{\omega}^u(s) \geq V_{\omega}^{u'}(s)$. That is, Z is a partition of Ω such that each subset forming the partition consists of a set of parameters that all agree on the ordering of the set of decisions according to value. Such a partition exists since the number of possible decision orderings is finite. Starting from the definition of EVOI, we have:

$$\begin{aligned} & EVOI(q, \psi) + V_{\psi}^* \\ &= \sum_{j=1}^k \Pr(q = j) V_{\psi|q=j}^* = \sum_{j=1}^k \Pr(q = j) V_{\psi|q=j}^{u_{\psi|q=j}^*} \\ &= \sum_{j=1}^k \sum_{\zeta \in Z} \Pr(q = j) \Pr(\omega \in \zeta | q = j) V_{\psi|q=j, \omega \in \zeta}^{u_{\psi|q=j}^*} \\ &= \sum_{j=1}^k \sum_{\zeta \in Z} \Pr(\omega \in \zeta) \Pr(q = j | \omega \in \zeta) V_{\psi|q=j, \omega \in \zeta}^{u_{\psi|q=j}^*} \\ &\leq \sum_{\zeta \in Z} \Pr(\omega \in \zeta) \max_{j'} \sum_{j=1}^k \Pr(q = j | \omega \in \zeta) V_{\psi|q=j, \omega \in \zeta}^{u_{\psi|q=j'}^*} \\ &= \sum_{\zeta \in Z} \Pr(\omega \in \zeta) \max_{j'} V_{\psi|w \in \zeta}^{u_{\psi|q=j'}^*}. \end{aligned}$$

Now let $\zeta_j = \{\zeta \in Z : \forall \omega \in \zeta, \arg \max_{j'} V_{\omega}^{u_{\psi|q=j'}^*} = j\}$, i.e., ζ_j is the collection of all ζ that prescribe value at least

as high to $u_{\psi|q=j}^*$ as $u_{\psi|q=j'}^*$ for all j' . Then, continuing from above, we have:

$$\begin{aligned} & \sum_{\zeta \in Z} \Pr(\omega \in \zeta) \max_{j'} \left\{ V_{\psi|w \in \zeta}^{u_{\psi|q=j'}^*} \right\} \\ &= \sum_{\zeta_j} \Pr(\omega \in \zeta_j) \max_{j'} \left\{ V_{\psi|w \in \zeta_j}^{u_{\psi|q=j'}^*} \right\} \\ &\leq \sum_{\zeta_j} \Pr(\omega \in \zeta_j) V_{\psi|w \in \zeta_j}^* \\ &= EVOI(q', \psi) + V_{\psi}^*, \end{aligned}$$

where recall q' is the k -response decision query over the set $\{u_{\psi|q=j}^*\}_{j=1}^k$. \square

Thus, we have shown that the set of k -response decision queries is EVOI-sufficient – this means that when restricted to asking k -response queries, there is no loss in considering only k -response decision queries. Furthermore, embedded in the proof above is a procedure for constructing a k -response decision query that is at least as valuable as a given k -response query. Repeating this procedure iteratively, then, would converge to a k -response decision query q_k^* with the property that when j is the response, the j^{th} decision in the set queried by q_k^* is the new Bayes-optimal decision. As a consequence, any EVOI-optimal k -response decision query must have this property. We will revisit this point in Section 7 where we describe the algorithms for k -response decision query selection by Viappiani and Boutilier (2010).

6 Nonmyopic k -Response Query Selection

Here we turn to the nonmyopic setting and ask: what query trees of depth- n should the agent consider when each query is constrained to having k responses? We begin our analysis by defining EVOI for query trees. Then we show that the set of depth- n k -response *decision-set* query trees contains a query tree with EVOI at least as high as any other depth- n k -response query tree. Lastly, we show that the same is *not* true for depth- n k -response decision query trees; i.e., Theorem 1 does not generalize to nonmyopic query selection.

6.1 Expected Value of Information for Query Trees

Let $M_n(Q)$ represent the set of all depth- n query trees that select queries only from query set Q . When a depth- n query tree μ is used to select n queries, the result is a trajectory of queries and responses. Let $X(\mu)$ denote the random variable where $\Pr(X(\mu) = j)$ represents the probability that the j^{th} such trajectory is realized when μ is used to select n queries.

Similarly to asking a single query, using μ to select n queries results in an updated posterior at the leaves; let $\psi|X(\mu) = j$ refer to the posterior distribution induced

when the j^{th} query and response trajectory is realized, which leads to a possibly new Bayes-optimal decision $V_{\psi|X(\mu)=j}^*$. Thus, we can write the EVOI associated with μ under ψ as

$$\text{EVOI}(\mu, \psi) = V_{\psi|X(\mu)}^* - V_{\psi}^*,$$

exposing the intuitive fact that the EVOI of any depth- n k -response query tree can be thought of as the EVOI of an equivalent k^n -response query, as stated by the following lemma:

Lemma 1. (Depth- n k -response query trees can be represented as k^n -response queries.)

For all finite decision problems and for all uncertainty ψ over them, for all depth- n k -response query trees μ , there exists a k^n -response query q^μ such that the EVOI of μ is equal to the EVOI of q^μ , implying that

$$\sup_{\mu \in M_n(Q_k)} \{\text{EVOI}(\mu, \psi)\} \leq \sup_{q \in Q_{k^n}} \{\text{EVOI}(q, \psi)\}.$$

Proof. Consider an arbitrary depth- n k -response query tree μ . Since there are k^n possible trajectories of queries and responses resulting from using μ to select a trajectory of n queries, we can construct a k^n -response query q^μ such that $\Pr(q^\mu = j|\omega) = \Pr(X(\mu) = j|\omega)$. Thus, $\text{EVOI}(\mu, \psi) = \text{EVOI}(q^\mu, \psi)$ since μ and q^μ are interchangeable in terms of their effect on ψ . \square

We will say that a depth- n k -response query tree set M is *Depth- n k -Response EVOI-Sufficient* if M always contains a query tree with EVOI at least as high as any other depth- n k -response query tree (hereafter, we will omit the dependence on k and n and refer to such query tree sets as EVOI-Sufficient).

Next we show that set of depth- n k -response decision-set query trees is EVOI-Sufficient, while the set of depth- n k -response decision query trees is *not* EVOI-Sufficient.

6.2 Decision Queries and Decision-set Queries in Nonmyopic Query Selection

First we show that the set of depth- n k -response decision-set query trees is EVOI-Sufficient. The following lemma provides a key step in proving this fact, reversing the relationship shown by Lemma 1 for the special case of depth- n k -response decision-set query trees and k^n -response decision queries:

Lemma 2. (k^n -Response decision queries can be represented as depth- n k -response decision-set query trees.)

For all finite decision problems and for all uncertainty ψ over them, for all k^n -response decision queries q , there exists a depth- n k -response decision-set query tree μ^q such that the EVOI of q is equal to the EVOI of μ^q , implying that

$$\max_{q \in D_{k^n}} \{\text{EVOI}(q, \psi)\} \leq \max_{\mu \in M_n(H_k)} \{\text{EVOI}(\mu, \psi)\}.$$




Figure 1: Diagram illustrating the main steps used to prove Theorem 2. Each arrow represents a statement that, for any query/query tree contained in the set at the tail of the arrow, a query/query tree with equal or higher EVOI must exist in the set at the head of the arrow. The three solid arrows, together with the fact that $M(H_k) \subset M(Q_k)$, imply Theorem 2 (represented by the dotted arrow).

Proof. Let q denote any k^n -response decision query. We prove the result by showing that it is always possible to construct a depth- n k -response decision-set query tree μ^q so that $\text{EVOI}(\mu^q, \psi) = \text{EVOI}(q, \psi)$.

We construct μ^q as follows. Let $Z_k(S)$ be any function that partitions a set S into k disjoint sets (where $|S|$ divisible by k), and let U^q denote the set of decisions queried by q . Then, construct μ^q such that $\mu^q(\psi)$ is the query that asks the decision-set query about sets $Z_k(U^q)$, and let $\mu^q(\psi|\mu^q(\psi) = j)$ be the query that asks the decision-set query about sets $Z_k(Z_k(U^q)_j)$, and so on. When taken together, responses to queries selected when using μ^q to select a trajectory of n k -response decision-set queries exactly determine which decision out of the original set U^q has maximum value for every possible ω .

Thus, if, under a particular ω , invoking μ^q to select a trajectory of n k -response decision-set queries determines that decision u_j^q is best out of the set, the response to q under ω would be j , and vice-versa; thus, for all ω , $\Pr(X(\mu^q) = j|\omega) = \Pr(q = j|\omega)$, which implies that $\text{EVOI}(\mu^q, \psi) = \text{EVOI}(q, \psi)$. \square

As a side note, the above construction implies that the set of k -response decision-set queries can be restricted to those containing decision sets of size k^{n-1} or less:

Corollary 1. For all finite decision problems and for all uncertainty ψ over them, the EVOI-optimal depth- n k -response decision-set query tree can be constructed using decision-set queries whose decision sets contain no more than k^{n-1} decisions.

We now put together Lemma 1, Lemma 2, and Theorem 1 to prove that the EVOI-optimal depth- n k -response query tree has EVOI no higher than the EVOI-optimal depth- n k -response decision-set query tree (Figure 1 for a visualization of the proof below):

Theorem 2. *(The set of depth- n k -response decision-set query trees is EVOI-sufficient.)*

For all finite decision problems and for all uncertainty ψ over them, the EVOI-optimal depth- n k -response query tree has EVOI equal to the EVOI of the EVOI-optimal depth- n k -response decision-set query tree:

$$\sup_{\mu \in M_n(Q_k)} \{EVOI(\mu, \psi)\} = \max_{\mu' \in M_n(H_k)} \{EVOI(\mu', \psi)\}.$$

Proof. Invoking Lemma 1, Theorem 1, and Lemma 2 sequentially, we have

$$\begin{aligned} \sup_{\mu \in M_n(Q_k)} \{EVOI(\mu, \psi)\} &\leq \sup_{q \in Q_{k^n}} \{EVOI(q, \psi)\} \\ &= \max_{q' \in D_{k^n}} \{EVOI(q', \psi)\} \\ &\leq \max_{\mu' \in M_n(H_k)} \{EVOI(\mu', \psi)\}, \end{aligned}$$

implying that

$$\sup_{\mu \in M_n(Q_k)} \{EVOI(\mu, \psi)\} = \max_{\mu' \in M_n(H_k)} \{EVOI(\mu', \psi)\},$$

since $M_n(H_k) \subseteq M_n(Q_k)$. \square

We now show that the set of depth- n k -response decision query trees is not EVOI-Sufficient, by constructing an example where no depth-2 binary-response decision query tree can have EVOI as high as the EVOI-optimal depth-2 binary-response decision-set query tree:

Theorem 3. *(The set of depth- n k -response decision query trees is not EVOI-Sufficient.)*

There exist finite decision problems under uncertainty ψ where the optimal depth- n k -response decision query tree has lower EVOI than the optimal depth- n k -response query tree, i.e., where

$$\sup_{\mu \in M_n(Q_k)} \{EVOI(\mu, \psi)\} > \max_{\mu' \in M_n(D_k)} \{EVOI(\mu', \psi)\}.$$

Proof. Consider a decision problem with four possible decisions u_1, u_2, u_3 and u_4 , where all 4! orderings over the values of each decision are supported by ψ . Lemma 1 and Lemma 2 together imply that the EVOI-optimal depth-2 binary-response query tree has EVOI equal to the EVOI of the myopically optimal 4-response decision query q^* . In this problem, asking q^* would allow the agent to act

optimally as a function of the response to q^* since there are only four possible decisions, which is achievable by the EVOI-optimal depth-2 binary-response query tree by Lemma 2. However, no depth-2 binary-response decision query tree exists that can meet this requirement.

To see this, note that the execution of a depth-2 binary-response decision query tree can eliminate at most two of the four decisions since all possible orderings of the decision values are supported by ψ – hence, no depth-2 binary-response decision query tree guarantees that the agent can act optimally after it selects 2 queries. \square

7 Algorithms and Computational Complexity

In Section 5 we showed that in the myopic setting, the set of all k -response queries can be replaced by the set of k -response decision queries at no loss, and in Section 6 we showed that in our nonmyopic setting, the set of all depth- n k -response query trees can be replaced by the set of depth- n k -response decision-set query trees at no loss. We conclude our analysis by discussing algorithms for myopic decision query selection and nonmyopic decision-set query selection, respectively.

7.1 Myopic k -Response Query Selection

We showed through Theorem 1 that computing the EVOI-optimal k -response query can be reduced to computing the EVOI-optimal k -response decision query. In fact, the problem of selecting the EVOI-optimal k -response decision query has been studied in recent work by Viappiani and Boutilier (2010) (our decision queries correspond to their “noiseless choice queries”).

First, they show that an EVOI-optimal k -response decision query has the property that when decision u is the response, u is the new Bayes-optimal decision (in Section 5 we showed that this fact can be understood as a of Theorem 1). Further, they prove that EVOI is *submodular* in k for k -response decision queries satisfying this property, implying that a greedily constructed k -response decision query has EVOI provably close to that of the EVOI-optimal k -response decision query (Nemhauser et al., 1978). Below we summarize their analysis of this greedy algorithm compared to a naive exhaustive algorithm.¹

Let $m = |U|$, and let the computational complexity of executing a single Bayes update (Equation 1) be $\mathcal{O}(l)$,

¹We omit the description of a third iterative algorithm presented by Viappiani and Boutilier (2010) (which can be understood as repeatedly applying the query improvement procedure embedded in the proof of Theorem 1, but specialized to decision queries), as its theoretical properties are less understood; however, this does not preclude its potential as an effective heuristic in practice.

where l is a measure of the size of the problem (which we leave undefined here because we will simply count how many such updates are performed by the different algorithms).

Exhaustive k -Response Decision Query Selection Algorithm. Exhaustively evaluate each possible k -response decision query and select the best one, which has computational complexity $\mathcal{O}(m^k kl)$.

Greedy k -Response Decision Query Selection Algorithm. Approximate the EVOI-optimal k -response decision query by a greedy procedure, at each step adding the decision contributing maximum EVOI. This algorithm enjoys the guarantee that the EVOI of the k -response decision query constructed is within a factor of $1 - (\frac{k-1}{k})^k$ (at worst 63%) of EVOI-optimal, and has computational complexity $\mathcal{O}(k^2 ml)$.

7.2 Nonmyopic k -Response Query Selection

Although Viappiani and Boutilier (2010) do not discuss algorithms for nonmyopic query selection, we can make use of the same algorithms as described above in the nonmyopic setting by exploiting our theoretical results.

Namely, combining Theorem 1 with Theorem 2 implies that we can compute the EVOI-optimal depth- n k -response query tree through two steps: (1) compute the EVOI-optimal k^n -response decision query q^* ; (2) construct a depth- n k -response decision-set query tree μ^* yielding the same EVOI as q^* .

Working backwards, step (2) can be implemented by the procedure described in our proof of Lemma 2, which involves computing any size- k partition of the k^n decisions queried by q^* for all k^n nodes of the tree. Since each of these k^n computations is $\mathcal{O}(k^n)$, step (2) has complexity $\mathcal{O}(k^{2n})$. Since step (1) can be implemented by either the exhaustive or the greedy algorithm above, we obtain two algorithms for depth- n k -response decision-set query tree selection:

Exhaustive Depth- n k -Response Decision-set Query Tree Selection Algorithm. This algorithm implements step (1) using the exhaustive k -response decision query selection algorithm above, and so its computational complexity is $\mathcal{O}(m^{k^n} k^{nl})$.

Greedy Depth- n k -Response Decision-set Query Tree Selection Algorithm. This algorithm approximates step (1) using the greedy k -response decision query selection algorithm described above, and so it has computational complexity $\mathcal{O}(k^{2n} ml)$ while offering the guarantee that the EVOI of the query tree computed is within a factor of $1 - (\frac{k^n-1}{k^n})^{k^n}$ (again, at worst 63%) of EVOI-optimal.

Thus we have shown the new result that the computational problem of selecting an EVOI-optimal depth- n k -response query tree can be reduced to selecting an EVOI-optimal k^n -response decision query.

8 Discussion

In this paper, we posed the question of what query set should be designed for an uncertain decision-making agent to allow the agent to select queries maximizing Expected Value of Information (EVOI), when the only restriction on what queries can be asked is that they must have exactly k possible responses (for some $k \geq 2$). In our myopic setting, we proved that the set of k -response decision queries is EVOI-Sufficient, which intuitively means that there is no benefit in considering additional k -response queries beyond decision queries. In our nonmyopic setting, where queries are used to construct depth- n query trees, we showed that the set of depth- n k -response decision query trees is *not* EVOI-Sufficient, but that the more general set of depth- n k -response decision-*set* query trees is in fact EVOI-Sufficient.

We then discussed algorithms developed in related work that can be directly applied to provably approximate k -response decision query selection, which exploit the submodularity of EVOI-optimal k -response decision query EVOI. Finally, we showed that the same algorithms apply to selecting depth- n k -response decision-set query trees by reducing depth- n k -response decision-set query tree selection to k^n -response decision query selection.

We note that this paper emphasized designing EVOI-sufficient k -response-query sets without regard to how humans may understand and answer queries from these sets. In some application domains, this may indeed be a practical challenge – in particular, it is clear that decision-set queries would be difficult for humans to answer unless the component decision-sets were to correspond to well-understood (by humans) categories of decisions. Studying how these types of queries can best be approximately conveyed in practical applications and how to take into account the cognitive burden they impose when evaluating them are important directions for future work.

Acknowledgments. We thank Alex Kulesza for numerous insightful conversations, and are grateful to the anonymous reviewers for their helpful feedback. Satinder Singh was funded by NSF grant IIS-1148668, and Edmund Durfee by NSF grant IIS-0964512. Any opinions, findings, conclusions, or recommendations expressed here are those of the authors and do not necessarily reflect the views of the sponsors.

References

- Paolo Viappiani and Craig Boutilier. Optimal Bayesian recommendation sets and myopically optimal choice query sets. In *Proceedings of the Annual Conference on Neural Information Processing Systems (NIPS)*, pages 2352–2360, 2010.
- Kathryn Chaloner and Isabella Verdinelli. Bayesian experimental design: A review. *Statistical Science*, 10:273–304, 1995.
- Burr Settles. Active learning literature survey. Technical Report 1648, University of Wisconsin – Madison, 2009.
- Robert D. Nowak. The geometry of generalized binary search. *IEEE Transactions on Information Theory*, 57(12):7893–7906, 2011.
- Darius Braziunas and Craig Boutilier. Elicitation of factored utilities. *AI Magazine*, 29(4):79–92, 2008.
- Kevin Regan and Craig Boutilier. Eliciting additive reward functions for Markov decision processes. In *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence (IJCAI)*, pages 2159–2164, 2009.
- Maya Cakmak, Crystal Chao, and Andrea L. Thomaz. Designing interactions for robot active learners. *IEEE Transactions on Autonomous Mental Development*, 2010.
- Warren B. Powell and Peter Frazier. *Optimal Learning*. John Wiley and Sons, 2008.
- 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)*, pages 363–369, 2000.
- Robert Cohn, Michael Maxim, Edmund Durfee, and Satinder Singh. Selecting operator queries using expected myopic gain. In *Proceedings of the International Conference on Intelligent Agent Technology (IAT)*, 2010.
- Darius Braziunas and Craig Boutilier. Local utility elicitation in GAI models. In *Proceedings of the Twenty-First Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 42–49, 2005.
- Craig Boutilier. A POMDP formulation of preference elicitation problems. In *Proceedings of the Eighteenth National Conference on Artificial Intelligence (AAAI)*, pages 239–246, 2002.
- Finale Doshi, Joelle Pineau, and Nicholas Roy. Reinforcement learning with limited reinforcement: Using Bayes risk for active learning in POMDPs. *Artificial Intelligence*, 187-188:115–132, 2012.
- Søren L. Dittmer and Finn V. Jensen. Myopic value of information in influence diagrams. In *Proceedings of the Thirteenth Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 142–149, 1997.
- Valentina Bayer-Zubek. Learning diagnostic policies from examples by systematic search. In *Proceedings of the Twentieth Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 27–34, 2004.
- Robert Cohn, Edmund Durfee, and Satinder Singh. Comparing action-query strategies in semi-autonomous agents. In *Proceedings of the Twenty-Fifth Conference on Artificial Intelligence (AAAI)*, pages 1102–1107, 2011.
- Daniel Golovin and Andreas Krause. Adaptive submodularity: Theory and applications in active learning and stochastic optimization. *Journal of Artificial Intelligence Research (JAIR)*, 42(1):427–486, 2011.
- Aaron Wilson, Alan Fern, and Prasad Tadepalli. A Bayesian approach for policy learning from trajectory preference queries. In *Proceedings of the Twenty-Sixth Annual Conference on Neural Information Processing Systems (NIPS)*, pages 1142–1150, 2012.
- Susan Athey and Jonathan Levin. The value of information in monotone decision problems. Working paper, Stanford University, 2001.
- Sonia Chernova and Manuela Veloso. Interactive policy learning through confidence-based autonomy. *Journal of Artificial Intelligence Research (JAIR)*, 34(1):1–25, 2009.
- Kshitij Judah, Alan Fern, and Thomas G. Dietterich. Active imitation learning via reduction to i.i.d. active learning. In *Proceedings of the Twenty-Eighth Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 428–437, 2012.
- G.L. Nemhauser, L.A. Wolsey, and M.L. Fisher. An analysis of approximations for maximizing submodular set functions–I. *Mathematical Programming*, 14(1):265–294, 1978.