Journal of Artificial Intelligence Research 58 (2017) 777-796 Submitted 8/16; published 4/17

Some Properties of Batch Value of Information in the Selection Problem

Shahaf S. Shperberg SHPERBSH@POST.BGU.AC.IL
Solomon Eyal Shimony SHIMONY@CS.BGU.AC.IL
Dept. of Computer Science
Ben-Gurion University of the Negev
P.O. Box 653, Beer-Sheva 84105, Israel

Abstract

Given a set of items of unknown utility, we need to select one with a utility as high as possible ("the selection problem"). Measurements (possibly noisy) of item values prior to selection are allowed, at a known cost. The goal is to optimize the overall sequential decision process of measurements and selection.

Value of information (VOI) is a well-known scheme for selecting measurements, but the intractability of the problem typically leads to using myopic VOI estimates. Other schemes have also been proposed, some with approximation guarantees, based on submodularity criteria. However, it was observed that the VOI is not submodular in general. In this paper we examine theoretical properties of VOI for the selection problem, and identify cases of submodularity and supermodularity. We suggest how to use these properties to compute approximately optimal measurement batch policies, with an example based on a "wine selection problem".

1. Introduction

Decision-making under uncertainty is a domain with numerous important applications. Since these problems are intractable in general, special cases are of interest. In this paper, we examine the selection problem: given a set of items of unknown utility (but drawn from a known distribution), we need to select an item with as high a utility as possible. Measurements (possibly noisy) of item values prior to selection are allowed, at a known cost. The problem is to optimize the overall decision process of measurement and selection. Even with the severe restrictions imposed by the above setting, this decision problem is intractable (Tolpin & Shimony, 2012; Radovilsky, Shattah, & Shimony, 2006; Radovilsky & Shimony, 2008); and yet it is important to be able to solve, at least approximately, as it has numerous potential applications. This paper analyzes cases where the value of information (VOI) is submodular, which is a sufficient condition for achieving good approximate solutions to the selection problem with appropriate greedy algorithms.

Settings where the selection problem is applicable are in meta-reasoning, i.e. considering which time-consuming deliberation steps to perform before selecting an action (Russell & Wefald, 1991b, 1991a; Hay, Russell, Tolpin, & Shimony, 2012), as well as settings where the items to be selected are physical objects. Examples of the latter type are: oil exploration, locating a point of high temperature using a limited number of measurements (Krause & Guestrin, 2009), performing costly measurements in order to find the best time to hit a target when the system is modeled using a stochastic system estimator, such as a Kalman filter, and a good set of parameters for setting up an industrial imaging system (Tolpin & Shimony, 2012). The selection problem can also be seen as a special case of Bayesian optimization, and can be used to select (from a large batch of candidates) experiments to be performed (Azimi, Fern, & Fern, 2016).

In most of the above applications the item values are naturally represented as being dependent. A potential application used in this paper as a running example (see Section 3) is selecting one from a set of wine cases that have uncertain qualities. Similar applications are selecting a batch of applicants with unrelated backgrounds to interview (from a larger set of job applicants) before making a hiring decision, and selecting a set of apartments from disparate locations to visit (among the available listings) before making a rental or purchasing decision. In the latter type of problems, the item distributions are either truly independent, or independent in a practical sense: it is not possible or worthwhile to obtain statistics beyond individual marginal distributions. (Note that such independence assumptions may be unjustified with job applicants that have a similar background, or apartments in the same building.) In fact, in meta-reasoning in search, an equivalent independence assumption called "subtree independence" is commonly made (Russell & Wefald, 1991b), even though it does not truly hold in the underlying search domains. Likewise, for oil exploration one could make such an independence assumption as a reasonable first-order approximation if one is considering as items a set of disjoint oil fields that are not physically near each other.

The selection problem is also called a Bayesian ranking and selection problem (Raiffa & Schlaifer, 2000; Frazier, 2012; Swisher, Jacobson, & Yücesan, 2003), where in some works the measurements are assumed to be noisy samples of the utility value of the items (Frazier, 2012). A widely adopted scheme for selecting measurements (also called sensing actions in some contexts, or deliberation steps in the context of meta-reasoning) is based on value of information (VOI) (Russell & Wefald, 1991b, 1991a). Optimizing value of information is intractable in general, thus both researchers and practitioners often use various forms of myopic VOI estimates (Russell & Wefald, 1991b, 1991a; Hay et al., 2012) coupled with greedy search. The properties of VOI have been of interest to the research community for quite a while (Raiffa & Schlaifer, 2000).

Nonconcavity in the value of information, a notion akin to non-submodularity in a continuous setting, was examined in several papers (Radner & Stiglitz, 1984; Lara & Gilotte, 2007; Chadeand & Schlee, 2001). Submodularity is an important property, because in cases where the VOI is submodular, simple, greedy algorithms result in provably near-optimal policies (Krause & Guestrin, 2009, 2011; Papachristoudis & Fisher III, 2012). However, the VOI is not submodular in general (Radner & Stiglitz, 1984; Lara & Gilotte, 2007; Chadeand & Schlee, 2001; Krause & Guestrin, 2009), and in particular submodularity does not hold in the selection problem (Tolpin & Shimony, 2012), even in a very limited case involving only two items. Other types of approximate solutions for the batch version of the selection problem exist (Reches, Gal, & Kraus, 2013), with theoretical bounds on the approximation error not based on submodularity. The latter paper also proved that the selection problem is NP-hard.

Specifically, the selection problem analyzed in this paper is as follows (see Section 2 for the formal definition). We have a set of items $\mathcal{I}$, each of which has some unknown value (or utility). The utility of each item is a random variable, and the joint distribution over the utilities of the items is known. It is possible to perform measurements on an item, thereby obtaining information about its utility. Measurements have a cost, specified by a known cost function $\mathcal{C}$, which is usually an additive cost function. After performing measurements, the decision-maker selects one item. We assume a risk-neutral decision-maker, and thus the decision maker always selects an item that has the highest expected utility given the observations. The problem is to find a policy of performing measurements such that the utility of the selected item minus the cost of measurements has a maximum expected value. In some settings (Tolpin & Shimony, 2012; Azimi et al., 2016), a *measurement budget* is also specified, and a policy is considered valid only if this budget is not exceeded. Some *budgeted* applications (Azimi et al., 2016) optimize just the expected value of the selection (not factoring in the measurement costs), but subject to the measurement budget constraint. Another common constraint is requiring that only a measured item may be selected, in which case submodularity holds under quite general conditions (Azimi et al., 2016); see Section 2 for a discussion on how their results relate to this paper. The latter constraint is natural for risk-averse decision makers, e.g. in the above hiring decision application, we may not wish to take the risk of hiring anyone we have not interviewed.

There are two common selection problem settings: batch, and online (also called sequential, or conditional). In the online setting, the decision-maker performs some measurements, then based on the resulting observations can decide on whether or not to perform additional measurements, etc. A policy in this case is essentially a conditional plan. In the batch setting, the decision-maker decides on a set (batch) of measurements to perform. The measurements are done essentially "in parallel", in the sense that the decision-maker does not get to perform additional measurements after receiving observations from previous ones. Given the observed results, the decision-maker then needs to make the final selection of an item. In this paper we consider only the batch setting.

In the batch setting, the value of information is the expected value of the best item given the observations, minus the expected value of the best item according to the initial (prior) distribution. That is, before receiving the information, there is some item that has the best expected value, which we call the "current best" item α. For simplicity we assume that this item is unique. After receiving the observations O, some other item $\beta(O)$ may have the highest expected value. The expected value of the difference $u_{\beta(O)} - u_\alpha$ is the value of information (VOI). Note that both the identity of the resulting best item β and its utility depend on O. In the batch setting with perfect observations, the distribution over the observed values is equal to the utility distribution of the respective item, and a set of measurments is fully specified by a set of items to be measured. Thus finding an optimal policy can be done by finding a set of measurements S to perform that has the highest expected VOI minus cost (also called the *net VOI*). In this paper, we consider mostly the case of perfect observations, i.e. where as a result of performing a measurement on an item, its precise utility value becomes known. In general, measurements can generate noisy (imperfect) observations. We briefly point out the cases where our results can be extended beyond perfect observations.

The theoretical results in this paper (Section 2) are as follows: the expected value of perfect information in the batch setting is neither submodular nor supermodular, in general. However, submodularity holds in the following cases:

- Theorem 1: the item utilities are jointly independent, and the utility of the currently best item α is known.
- Theorem 3: the utility of α is known, and there are only at most two additional items (may be dependent).
- Theorem 4: the item utilities are jointly independent, the utility of α is sufficiently high (condition C1), and the decision maker is constrained to *always* measure α.

Theorem 1 is important because even this simple setting leads to an NP-hard selection problem (Theorem 2). We also show by providing simple discrete-valued counterexamples where attempting to generalize these theorems fails. Theorem 3 cannot be generalized to more items, as a counterexample with three items (in addition to α) is presented. In Theorem 4, violating condition C1 makes the theorem break, even with two items (in addition to α). Finally, we capitalize on the submodularity results (Theorems 1 and 4) by suggesting a simple "compound" greedy scheme in Section 3 for near-optimal solution of the selection problem, and compare its performance to the standard greedy algorithms on a wine quality dataset.

2. Main Results

We begin by formally defining the perfect information batch selection problem.

Definition 1 (perfect information batch selection setting)**.** *Let $\mathcal{I} = \{I_0, I_1, ..., I_n\}$ be a set of $n+1$ items with uncertain utility, represented by r.v.s $X_0, ..., X_n$. We assume w.l.o.g. that the current best item α is item I_0. For a cost C_i we can measure I_i, obtaining a (perfect) observation of the utility of this item. We select a subset $S \subseteq \mathcal{I}$ to be measured as a batch, for a total cost of $\sum_{I_i \in S} C_i$, after which we observe the results O (the true utilities of the items in S), and select a final item $I_f(O)$ that has the highest expected utility given the observations.*

The *optimization version* of the perfect information batch selection problem is: under the perfect information batch selection setting, find the set that achieves:

$$\max_{S \in \mathcal{I}} (E_S[I_f(O)] - \sum_{I_i \in S} C_i) \tag{1}$$

where the subscript of E indicates the set of random variables over which the expectation is performed, and O are the observations due to measuring the items in S. Optionally, in the *budget limited* version of the selection problem, we are given a budget limit C and need to optimize S under the additional constraint:

$$\sum_{I_i \in S} C_i \leq C \tag{2}$$

In order to simplify some of the proofs below, we make the additional assumption that the decision maker always picks item α if its expected utility is at least as high as that of all other items, given the past observations. That is, "among equals, prefer item α". Denoting the expected value of X_i by μ_i, note that by construction $\mu_\alpha \geq \mu_i$ for all $i > 0$. For a set of items $S \subseteq \mathcal{I}$, denote the expected value of information of a (perfect) observation of the utility of all these items by $VPI(S)$, defined as the expected value $E_S[I_f(O)] - \mu_\alpha$ (with expectation taken over all possible observations on S). Denote by p_i the PDF of random variable X_i.

Example 1. *Consider a wine selection problem with quality distributions similar to Figure 2. Suppose that one wine case α that we wish to purchase has a known quality of $u_\alpha = 8$. We have been offered two additional options, one with a quality distribution X_1 uniformly distributed in $\{5, 6, 7, 8, 9\}$, i.e. $\mu_1 = 7$, the other (X_2) with quality in $\{4, 10\}$, again uniformly distributed, so $\mu_2 = 7$. Suppose that our utility scale is linear in the quality, that all wines cases cost the same (or that cost has already been factored in negatively into the quality). Since the wines are not known to be related, we assume that the quality distributions are independent. Testing some of the wine cases is possible (at a known cost, though we ignore such costs at present), thereby revealing their true quality. Note that if we test no wines, then we should rationally pick the α wine for a quality of 8. If we choose to test wine case 2 prior to the purchase, then with probability 0.5 its quality is revealed as 10, and we purchase it, thereby gaining 2. Otherwise, stick with α, and gain nothing. On the average we gain 1, so $VPI(\{X_2\}) = 1$.*

2.1 Batch VOI when the Utility of the Currently Best Item is Known

Theorem 1. *For a perfect information batch selection setting with independent item utility distributions, where the utility u_α of the currently best item α is known, the value of perfect information $VPI(S)$ is a submodular set function.*

Proof: Due to independence, an item that has not been measured will never be selected (except for α). Measured item I_i will be selected if its utility is observed to be greater than u_α and the rest of the observed items. (For conciseness, we use $M(S)$ to denote $\max(\{u_i | I_i \in S\})$, and $M(X_S)$ to denote the respective random variable $\max(\{X_i | I_i \in S\})$.) Therefore, the VPI of observing (the utility of) a set of items S that does not include α is:

$$\begin{aligned} VPI(S) &= \int_{M(S) > u_\alpha} (M(S) - u_\alpha) \prod_{I_i \in S} p_i(u_i) du_i \\ &= \int (\max(M(S), u_\alpha)) - u_\alpha) \prod_{I_i \in S} p_i(u_i) du_i = E_S[\max(M(X_S), u_\alpha))] - u_\alpha \quad (3) \end{aligned}$$

Now write down the difference in VPI between $S \cup \{I\}$ and S, for some item $I \notin S$ using Equation 3:

$$\begin{aligned} VPI(S \cup \{I\}) - VPI(S) &= (E_{S \cup \{I\}}[\max(M(X_{S \cup \{I\}}), u_\alpha)] - u_\alpha) \\ &\quad -(E_S[\max(M(X_S), u_\alpha)] - u_\alpha) \\ &= E_{S \cup \{I\}}[\max(M(X_S), X_I, u_\alpha) - \max(M(X_S), u_\alpha)] \\ &= E_{S \cup \{I\}}[\max(X_I - \max(M(X_S), u_\alpha), 0)] \end{aligned}$$

Consider the difference in VPI for set $S' = S \cup \{J\}$ for some item $J \neq I$, $J \notin S$. We have:

$$\begin{aligned}
VPI(S' \cup \{I\}) - VPI(S') &= E_{S'_{\cup}\{I\}}[\max(X_I - \max(M(X_{S'}), u_\alpha), 0)] \\
&= E_{S_{\cup}\{I,J\}}[\max(X_I - \max(M(X_S), X_J, u_\alpha), 0)] \\
&\leq E_{S_{\cup}\{I,J\}}[\max(X_I - \max(M(X_S), u_\alpha)), 0)] \\
&= E_{S_{\cup}\{I\}}[\max(X_I - \max(M(X_S), u_\alpha)), 0)] \\
&= VPI(S \cup \{I\}) - VPI(S)
\end{aligned}$$

Where the inequality follows due to removing a term from the (negated) maximization. We have obtained that for all sets S that do not include α, the difference in VPI is non-increasing as S (setwise) increases. Therefore, $VPI(S)$ is a submodular function of S. □

Corollary 1. *Theorem 1 also holds given only a distribution over u_α, if there is no way to obtain additional information about u_α. That is because an optimal (risk neutral) decision maker would have to act as if $u_\alpha = \mu_\alpha$.*

A similar argument leads to a generalization to noisy observations: although Theorem 1 is stated in terms of perfect information, this is not an inherent limitation. Consider a more general setting where measurements are noisy, but the value of each item can be measured only once. In this setting, one can simply use the expected posterior value instead of the actual value when making the decision, and our results still apply. However, in settings where the measurement types on an item are allowed to vary (e.g. allow a choice between one and two conditionally independent measurements, or a choice of measurements that reveal different features of an item), it is well known that submodularity does not hold (Frazier & Powell, 2010; Tolpin & Shimony, 2012).

Observe that Theorem 1 is relevant to additional settings. First, consider the special case where $u_\alpha = 0$. In this case, action α can be re-cast as making no selection at all, and the conditions of the theorem hold if all items have a non-positive prior expected value. This is actually reasonable when items with uncertain value are being sold to our decision-making agent, as the seller wishes to gain from the sale, and presumably would not wish to sell an item for less than its expected value.

Another setting in which Theorem 1 applies is if the agent is *not allowed* to select an item unless its value has been measured or was known previously. In this setting the requirement that u_α is greater than the expectation of all the rest of the items can be dropped, and the results can be made significantly more general (Azimi et al., 2016), as shown in Lemma 1 in the latter paper. The lemma states that the expectation of the maximum of a set of random variables is monotonic non-decreasing and submodular, and does not even require that the variables be independent. Lemma 1 can thus also be used to prove our Theorem 1. It is interesting to note that in order to apply their Lemma 1, perfect observations were required (Azimi et al., 2016). However, unlike our Theorem 1, the lemma cannot be easily applied if we relax the perfect information limitation, as that would lead to an apparent contradiction due to dependencies, as discussed in Section 2.1.2.

2.1.1 Complexity of the Selection Problem

The batch measurement selection problem was shown to be NP-hard (Reches et al., 2013), in a setting where multiple noisy measurements per item are allowed. We show that the problem gives rise to an NP-hard decision problem even if the observations are perfect.

Definition 2 (perfect information budget-limited batch selection decision problem (PBSP))**.** *In the perfect information batch selection setting from Definition 1, is there a subset $S \subseteq \mathcal{I}$ which has a total measurement cost not greater than C, and such that the expected utility of the final item I_f selected after observing the utility of the items in S, is at least U?*

Theorem 2. *The PBSP is NP-hard.*

The proof appears in Appendix A, by reduction from Knapsack to a PBSP restricted to the case where u_α is known to be 0, and where the unknown item utility distributions are independent over $\{-1, 1\}$. It follows immediately from these restrictions that the perfect information budget-limited batch selection decision problem remains NP-hard under the conditions of Theorem 1. Note, however, that if we also restrict the measurement costs in PBSP to be all equal, a greedy algorithm would deliver an optimal subset, and thus a trivial polynomial time solution to the PBSP. Whether the PBSP with equal costs but with *arbitrary* independent discrete distributions is NP-hard is an open problem.

2.1.2 VPI in the Presence of Dependencies

We now consider the perfect information batch selection setting, without the independence assumption. With dependencies the amount of information obtained by additional observations, having already made some observations, is ususally reduced. Intuition would suggest that the same would therefore occur for the VPI as well. Indeed, for $n = 2$ the VPI is still subadditive. For example, the wine selection problem in example 1 with the same marginal distributions, but where X_1 and X_2 are dependent, falls under this case.

Theorem 3. *For a batch selection setting with 3 items, where the utility of the currently best item α is known, the value of perfect information is subadditive, i.e. $VPI(\{1\}) + VPI(\{2\}) \geq VPI(\{1, 2\})$.*

Proof: Note that, unlike the independent case, when observing only one item it is actually possible to select either the other, unobserved item, or α. This results in:

$$VPI(\{1\}) = \int_{\max(u_1, \mu_{2|1}(u_1)) > u_\alpha} (\max(u_1, \mu_{2|1}(u_1)) - u_\alpha) p_1(u_1) du_1$$

where $\mu_{2|1}(u_1)$ is the expected utility of item 2 given that the item 1 was observed to have utility u_1, defined as:

$$\mu_{2|1}(u_1) = \int_{u_2} p_2(u_2|u_1) u_2 du_2$$

the VPI for item 2 is defined symmetrically, exchanging the roles of items 1 and 2 in these equations.

The value of information for observing both item 1 and item 2 (followed by selecting the best of them or α, whichever has maximal utility) is:

$$VPI(\{1,2\}) = \int\int_{\max(u_1,u_2)>u_\alpha} (\max(u_1,u_2) - u_\alpha) p_{1,2}(u_1,u_2) du_1 du_2 \tag{4}$$

Separating out the domain we can write:

$$VPI(\{1,2\}) = \int_{u_1>u_\alpha}\int_{u_2\le u_1}(u_1-u_\alpha)p_{1,2}(u_1,u_2)du_1du_2 + \int_{u_2>u_\alpha}\int_{u_1<u_2}(u_2-u_\alpha)p_{1,2}(u_1,u_2)du_1du_2$$

Denote the first integral by J_1 and the second by J_2, for convenience. We now rewrite $VPI(\{1\})$ as a sum over the regions:

$$VPI(\{1\}) = \int_{u_1>u_\alpha}(\max(u_1,\mu_{2|1}(u_1)) - u_\alpha)p_1(u_1)du_1 + \int_{\{u_1|u_1\le u_\alpha \wedge \mu_{2|1}(u_1)>u_\alpha\}}(\mu_{2|1}(u_1) - u_\alpha)p_1(u_1)du_1$$

Now, dropping the $\mu_{2|1}(u_1)$ from the maximization in the first integral and noting that the second integral is non-negative, we get:

$$\begin{aligned}
VPI(\{1\}) &\ge \int_{u_1>u_\alpha}(u_1-u_\alpha)p_1(u_1)du_1 \\
&= \int_{u_1>u_\alpha}\int_{u_2\le u_1}(u_1-u_\alpha)p_{1,2}(u_1,u_2)du_1du_2 + \int_{u_1>u_\alpha}\int_{u_2>u_1}(u_1-u_\alpha)p_{1,2}(u_1,u_2)du_1du_2 \\
&= J_1 + \int_{u_1>u_\alpha}\int_{u_2>u_1}(u_1-u_\alpha)p_{1,2}(u_1,u_2)du_1du_2 \ge J_1
\end{aligned}$$

Likewise we show that $VPI(\{2\}) \ge J_2$, and thus $VPI(\{1\}) + VPI(\{2\}) \ge VPI(\{1,2\})$. □

Unfortunately, this submodularity result has no *practical* use, as it does not generalize to $n \ge 3$, as is evident from the following counterexample. Let $u_\alpha = 10$, and we have 3 additional items with utility distributed as binary variables, with values $\{L, H\}$. The dependency is "parity", that is, exactly an even number of the items have value H, and the rest have value L. The distribution over the 4 possible legal configurations is uniform, i.e. each has probability 0.25. The utility values are: $u_{1_L} = u_{2_L} = 5$, and $u_{1_H} = u_{2_H} = 13$, so that $\mu_1 = \mu_2 = 9 < u_\alpha$. For the 3rd item, we have: $u_{3_L} = 0$, and $u_{3_H} = 18$, so that $\mu_3 = 9 < u_\alpha$.

Note that the marginal distribution over each of the items is uniform, and remains uniform given the observation of one other item, i.e. the variables are pairwise independent. The individual VPIs are therefore:

$$VPI(\{1\}) = 0.5 \times 0 + 0.5 \times (13 - 10) = 1.5$$

and due to symmetry we also have $VPI(\{2\}) = 1.5$. Having observed both items 1 and 2, the utility of item 3 is known with certainty, and it is selected if known to have value H. Therefore we have:

$$VPI(\{1,2\}) = 0.25 \times (0 + (18-10) + (18-10) + (13-10)) = 4.75 > VPI(\{1\}) + VPI(\{2\})$$

Finally, note that Lemma 1 in the work of Azimi et. al. (2016) does not require independence, and can be used to show that VOI of perfect measurements is submodular if we do not allow the agent to select an unmeasured item. The above counterexample shows that the applicability of Lemma 1 cannot easily be extended to allow imprecise measurements. Allow a very noisy measurement for item 3, e.g. $P(ObserveHIGH|u_3 = H) = 0.51$, $P(ObserveHIGH|u_3 = L) = 0.49$. The VOI values change only slightly, but now the requirement that only measured items can be selected is met. However, as shown above, the value of information is not submodular.

2.2 Batch VOI when the Utility of the Currently Best Item is Unknown

Consider now that we are given a distribution over u_α, but unlike corollary 1, additional information about u_α can be obtained. For simplicity, consider just the case with 2 items. In the well-known case exhibited in figure 1, we can see that it is not possible, by observing only one item, to make an optimally behaving agent change the choice from α to β. So the individual VPI are zero. But since there is some non-zero probability that u_α is less than u_β, observing both items it is possible that β will be selected to increase the utility, therefore we have $VPI(\{\alpha, \beta\}) > 0$. In this case the value of perfect information is supermodular, but does this hold in general for 2 items?

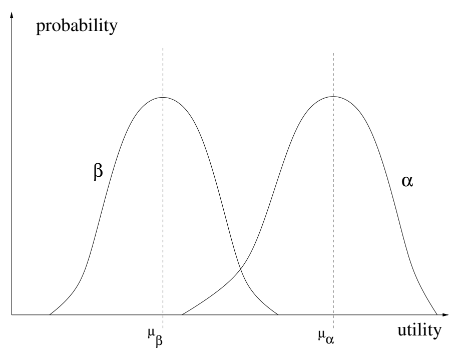

Figure 1: Utility distributions with supermodular VPI

Consider the distributions: u_α is evenly distributed: $P(u_\alpha\!=\!0) = P(u_\alpha\!=\!10) = 0.5$, and u_β is distributed as $P(u_\beta\!=\!1) = 0.7$, $P(u_\beta\!=\!11) = 0.3$. We get $(\mu_\alpha = 5) > (\mu_\beta = 4)$. The individual values of perfect information are:

$$VPI(\{\alpha\}) = 0.5 \times 0 \ + \ 0.5 \times (4 - 0) \ = \ 2$$

$$VPI(\{\beta\}) = 0.7 \times 0 \ + \ 0.3 \times (11 - 5) \ = \ 1.8$$

For observing both items, we get:

$$\begin{aligned} VPI(\{\alpha, \beta\}) &= 0.5\times0.3\times(11-10) + 0.5\times0.3\times(11-0) + 0.5\times0.7\times(1-0) + 0.5\times0.7\times0 \\ &= 2.15 \ < \ VPI(\{\alpha\}) + VPI(\{\beta\}) \end{aligned}$$

This example is discouraging, since we have neither submodularity nor supermodularity. An interesting question is about the VPI among sets of observations that *must* include an observation of the currently best item. In general, the VPI of such sets is neither submodular nor supermodular, as shown by the following counterexample.

We have 3 items, with distributions as follows. Current best item α, distributed: $P(u_\alpha = 20) = P(u_\alpha = 0) = 0.5$. Second best item β, distributed: $P(u_\beta = 9) = P(u_\beta = 5) = 0.5$, and third item γ, distributed $P(u_\gamma = 6) = P(u_\gamma = 2) = 0.5$. This gives us: $\mu_\alpha = 10$, $\mu_\beta = 7$ and $\mu_\gamma = 4$. If we observe only item α, if it has a low value we pick item β so we have:

$$VPI(\{\alpha\}) = P(u_\alpha = 0) \times (\mu_\beta - 0) = 3.5$$

If we also observe item β, this makes no difference as any possible utility value for item β is still higher than μ_γ. Likewise, observing α and γ, we still select α if $u_\alpha = 20$ and β if $u_\alpha = 0$. Therefore we have:

$$VPI(\{\alpha, \beta\}) = VPI(\{\alpha, \gamma\}) = VPI(\{\alpha\}) = 3.5$$

However, the value of observing all items is higher, since item γ may be better:

$$\begin{aligned} VPI(\{\alpha,\beta,\gamma\}) &= P(u_\alpha = 0) \times [P(u_\beta = 9) \times 9 + P(u_\beta = 5) \times (P(u_\gamma = 6) \times 6 + P(u_\gamma = 4) \times 5)] \\ &= 0.5 \times [0.5 \times 9 + 0.5 \times 5.5] \\ &= 3.625 > VPI(\{\alpha, \beta\}) + VPI(\{\alpha, \gamma\}) - VPI(\{\alpha\}) \end{aligned}$$

which clearly violates submodularity for sets containing observations of item α.

However, if u_α is always sufficiently high, i.e. if every possible value of u_α is no less than μ_β, submodularity does hold among such sets. In general, denote:

Condition C 1. $P(u_\alpha < \mu_i) = 0$ *for all items* i *other than* α.

Example 2. *Consider the same wine selection problem instance as in example 1, except that the quality of the* α *wine case is no longer known to be 8: instead its quality is uniformly distributed among* $\{7, 8, 9\}$. *With the qualities* X_1 *and* X_2 *being independent and distributed as in example 1, this example obeys condition C1.*

Formally, denote by $VPI^\alpha(S)$ the value of information of perfectly observing the utility of all items in S, as well as that of α. (This is equivalent to $VPI(S \cup \{\alpha\})$, but we wish to emphasize that this is a function of S not including α, hence the above notational variant.)

Theorem 4. *For a batch selection setting with jointly independent items where condition C1 holds, the value of perfect information* $VPI^\alpha(S)$ *is a submodular set function of* S.

Proof: Note that this theorem is a strict generalization of the corollary of Theorem 1, as condition C1 always holds trivially if α is the current best item and u_α is known. Condition C1 ensures that after the observations of items in S, the optimal policy must select either one of these observed items, or α. (This is not necessarily the case if C1 is violated.) Therefore the VPI here can be obtained in a manner similar to Equation 3, with integration over u_α:

$$\begin{aligned} VPI^\alpha(S) &= \int_{u_\alpha} \int_{M(S) > u_\alpha} (M(S) - u_\alpha) p_\alpha(u_\alpha) du_\alpha \prod_{I_i \in S} p_i(u_i) du_i \\ &= E_{S \cup \{\alpha\}}[\max(M(X_S), u_\alpha)) - u_\alpha] \end{aligned} \quad (5)$$

Similar to the proof of Theorem 1, write down the difference in VPI^{α} between $S \cup \{I\}$ and S, for some item $I \notin S$ using Equation 5:

$$\begin{aligned} VPI^{\alpha}(S \cup \{I\}) - VPI^{\alpha}(S) &= E_{S\cup\{I,\alpha\}}[\max(M(X_{S\cup\{I\}}), u_{\alpha})) - u_{\alpha}] \\ &\qquad - E_{S\cup\{\alpha\}}[\max(M(X_S), u_{\alpha})) - u_{\alpha}] \\ &= E_{S\cup\{I,\alpha\}}[\max(M(X_S), X_I, u_{\alpha}) - \max(M(X_S), u_{\alpha})] \\ &= E_{S\cup\{I,\alpha\}}[\max(X_I - \max(M(X_S), u_{\alpha}), 0)] \end{aligned}$$

Consider the difference in VPI for set $S' = S \cup \{J\}$ for some item $J \neq I$, $J \notin S$. We have:

$$\begin{aligned} VPI^{\alpha}(S' \cup \{I\}) - VPI^{\alpha}(S') &= E_{S'_{\cup}\{I,\alpha\}}[\max(X_I - \max(M(X_{S'}), u_{\alpha}), 0)] \\ &= E_{S_{\cup}\{I,J,\alpha\}}[\max(X_I - \max(M(X_S), X_J, u_{\alpha}), 0)] \\ &\leq E_{S_{\cup}\{I,J,\alpha\}}[\max(X_I - \max(M(X_S), u_{\alpha})), 0)] \\ &= E_{S_{\cup}\{I,\alpha\}}[\max(X_I - \max(M(X_S), u_{\alpha})), 0)] \\ &= VPI^{\alpha}(S \cup \{I\}) - VPI^{\alpha}(S) \end{aligned}$$

Therefore, $VPI^{\alpha}(S)$ is a submodular set function of S. □

3. Application of Results

A typical application of submodularity is in algorithms that compute near-optimal policies for selection in the perfect information batch selection setting. Consider for example a batch setting selection problem where the measurement cost function $\mathcal{C}$ is supermodular (or additive, as a special case common in applications). As a result, the optimal solution to the (batch setting) selection problem is to measure a set of items S that maximizes $VPI(S) - \mathcal{C}(S)$ (the net VPI), followed by selecting the item with the best expectation given the observations.

If we know the utility of item α, then $VPI(S) - \mathcal{C}(S)$ is submodular due to Theorem 1. We can thus use a standard greedy algorithm that starts with an empty candidate set S, and repeatedly adds to S items that have the highest net gain (best (marginal) VPI minus cost), until no item has a positive net gain. We call this method the **(additive) greedy** algorithm. According to a fundamental result by Nemhauser et. al. (1978), the greedy algorithm already guarantees an expected utility that is close to optimal for monotone submodular functions. The quality of the greedy algorithm in practice is usually much better than the guaranteed bounds, and a similar tendency can be seen in the example wine selection application below. The quality of the results seems to occur due to the fact that submodularity is a guarantee against "premature stopping" in the greedy algorithms[1]; deviations from optimality resulting from picking non-optimal items early-on seem much less problematic in practice than indicated by the worst case in theory (a factor of $1 - 1/e$ from optimal for the monotonic case, worse for non-monotonic which is the case here).

As cases where u_{α} is known may be rare, it is possible to use a similar scheme if u_{α} is not known, but the distributions obey condition C1. In this case, run the greedy algorithm

1. "Premature stopping" is a term used to mean that although there is a set of items with a combined VOI greater than its measurement cost, the (greedy) algorithm decides not to measure any additional items because their individual VOI is too low. See example 3 in Section 3.1.

twice: once for sets that do not contain measurements of item α, and once for sets that do contain such measurements; compare the expected value of both resulting measurement sets, and return the better of the two. We call this method the **compound greedy** algorithm. Again, Theorems 1 and 4 imply that the functions we optimize in both cases are submodular, thus the greedy algorithms return sets that are near-optimal.

3.1 Example Setting: Wine Selection

We examine an example application for the perfect information batch selection setting, and solve a selection problem on a typical set of items. A comparison of algorithm performance on such a dataset indicates the type of results one can observe with greedy optimization algorithms for the selection problem.

Definition 3 (Wine selection problem)**.** *Given a set of wine types $\mathcal{I} = \{I_0, ..., I_n\}$, each wine has an unknown quality, but a quality distribution is known for each type. In addition, for a known cost C_i, we can purchase and send a bottle of each wine type to a sommelier for analysis and quality determination (or taste it ourselves, for the few people who actually understand wine quality, though clearly not the authors of this paper). Which subset S of the wines (if any) needs to be sent to the sommelier in order to maximize the expected utility of testing and final decision? (That is, maximize the expected quality of the final selection, minus the sum of costs C_i of wines in S, i.e. the net VPI).*

The setting for the tests was based on the UCI white wine quality dataset (Cortez, 2009; Cortez, Cerdeira, Almeida, Matos, & Reis, 2009). The dataset contains over 1000 wines, with 11 feature values for each wine, such as pH, alcohol level, etc. The target attribute value (quality) is based on evaluations made by wine experts, and ranges from 1 (very bad) to 10 (excellent).

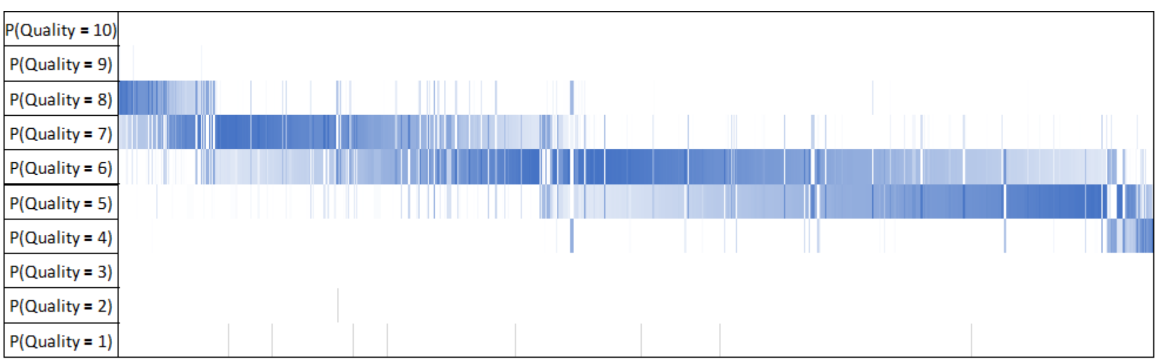

Figure 2: Wine quality distributions

Using this dataset, we constructed for each wine a quality distribution based on the quality distribution of all wines in the data set that had the same feature values. This distribution was adjusted by applying Kernel density estimation (KDE) using a Gaussian kernel and a rule of thumb (Silverman, 1986) that increases the kernel width as a function of the variance. This resulted in the wine quality distributions depicted in Figure 2, a distribution scatter plot where darker color indicates higher probability. Each value on the

X axis indicates a specific wine type, with wines sorted by expected quality value. The wine quality distributions are assumed to be independent.

Using the above distribution, the following experiments were conducted. Each experiment was on a set $\mathcal{I}$ of $n+1$ randomly picked wines from the dataset, where n was an experimental parameter, and for each wine a random cost C_i was drawn uniformly between 0.01 to 0.1 (assumed to be on the same scale as quality values). The wine with the best expected value from $\mathcal{I}$ is the α wine, the prior best. We then used 4 different methods to find the measurement policy (i.e. batch of wines to be tested).

1. **Exhaustive**: Every possible subset S of $\mathcal{I}$ (both with and without the alpha wine) was examined. The S which maximized the net VPI was returned. S here is the optimal (batch) measurement policy.

2. **Greedy (additive)** approach. The wines are kept sorted according to their myopic expected net VPI w.r.t. the current batch. A batch S is incrementally constructed, starting from the empty set: every iteration, the best candidate wine from $\mathcal{I} - S$ is added to S, as long as the net myopic VPI for adding this wine is positive.

3. **Greedy (rate)** approach. This greedy method is the same as the additive greedy approach, except that the wines are kept sorted according to their expected VPI divided by cost of the measurement. Once this value drops below 1 for all remaining wines, the algorithm returns the current batch. This approach is the same as in the paper by Azimi et. al. (2016), modified to the wine selection problem.

4. **Compound greedy** approach: Run the greedy (additive) algorithm twice: once for sets that do not contain the α wine, and once for sets that do contain α. Compare the expected net VPI of both resulting measurement sets, and return the better of the two.

For each of the following cases, we varied n, the number of items, from 1 to 20:

1. Known u_α, created by setting the value of u_α to its mean in each randomly picked item set. Here compound greedy is the same as simple greedy, so is not shown.

2. u_α unknown, but condition C1 holds (generated by random sampling of sets, rejecting sets where C1 did not hold).

3. Instances where condition C1 does not hold (and obviously unknown u_α).

The net VPI, averaged over 5 random item sets for each item set size, is shown in Figure 3.

Both standard greedy algorithms averaged 0.99 of the optimal net VPI, while compound greedy averaged slightly better at 0.993, all considerably better than the theoretical bound. It is interesting to observe that the greedy algorithms performed well even in many cases where the theorems do not guarantee submodularity, such as the cases where condition C1 did not hold (Figure 3 upper right). In some cases rate greedy performed better than both of the other methods, but a rate-based version of compound greedy (not shown) dominates rate-greedy. In extreme cases (which did not occur in the above runs) the net VPI is 0 for both rate and additive greedy, even though considerable net VPI is achievable. This occurs due to the "premature stopping" phenomenon caused by non-diminishing returns.

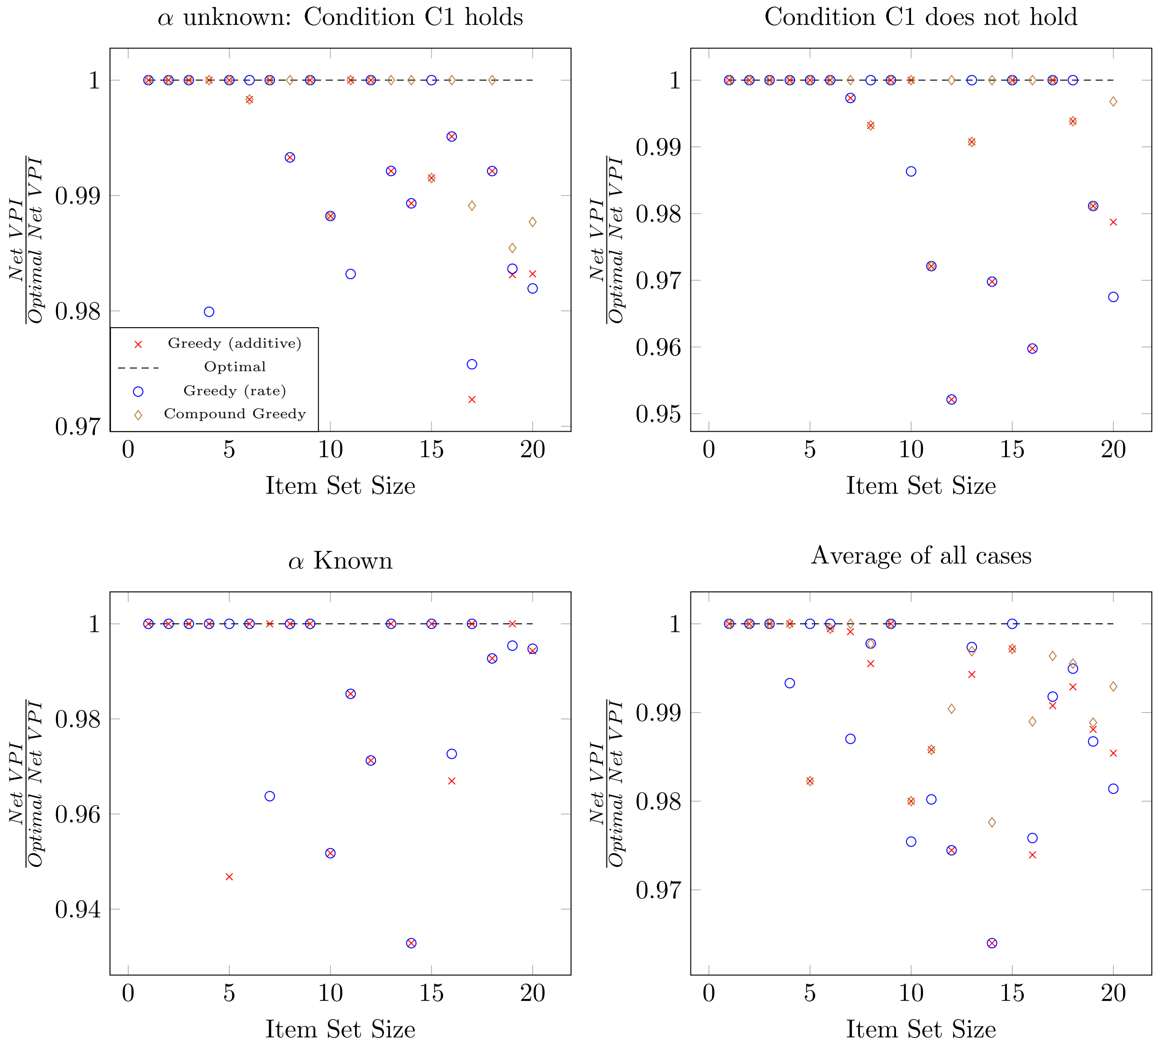

Figure 3: Comparison of net VPI for various item set sizes

Example 3. *Consider the case where condition C1 holds as in Example 2, with quality distribution of the α wine being uniform among $\{7, 8, 9\}$, but in addition the best possible quality in all the other items is no better than the $E(X_\alpha)$ value, such as when the only other choice is X_1 distributed uniformly among $\{6, 7, 8\}$. In this case the VPI of every singleton set is 0, similar to the situation depicted in Figure 1, whereas measuring both wines results in a gain of 1 with probability $\frac{1}{9}$, and thus $VPI^\alpha(\{X_1\}) = \frac{1}{9}$. This will cause the rate and additive greedy algorithms to incorrectly return an empty set of items to be measured. The compound greedy algrithm avoids exactly this pitfall.*

We now turn to the issue of computation time. All the above algorithms require evaluation of the VPI of a batch, which can itself be non-trivial. An initial naive implementation caused even the *greedy* algorithms to time out on sets of 20 wines. This sub-problem can be handled in the general case by approximating the VPI (i.e. expectation of the maximum) of each batch by sampling (Azimi et al., 2016). In the wine selection problem, however, we have independent discrete random variables with greatly overlapping domains, so we can cheaply compute the distribution of the maximum, and from there evaluate the expectation of the maximum exactly.

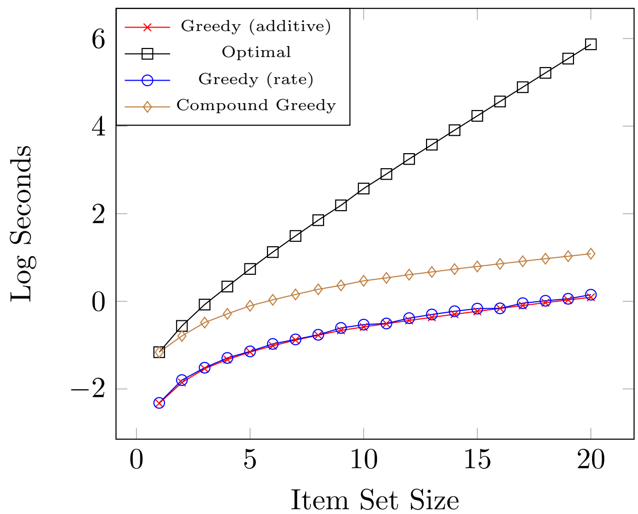

Figure 4: Algorithm runtime comparison: time vs. item set size

Runtimes for the algorithms appear in Figure 4, performed on an Intel(R) Core(TM) i7-4700HQ 2.40GHz with 8 GB RAM running Microsoft windows 8.1 x64, using multiple-thread implementations. The software was implemented in C# with optimizations. Clearly, the exhaustive method delivers the best net VPI, but its runtime is prohibitive for large sets of wines.

Both the additive and rate greedy were the fastest, with compound greedy roughly a constant factor slower. In fact, despite the improved VPI computation, this part still dominates the runtime, and adding caching of computations of random variable maximizations resulted in the compound greedy algorithm being only a few percent slower than the other greedy algorithms (not shown). Therefore, although the improvement due to the compound greedy algorithm appears small, it comes essentially for free and is thus worthwhile. The greedy algorithms appear to be scalable: an experimental run with $n = 100$ wines resulted in runtimes of approximately 200 seconds for each of the greedy algorithms (including compound greedy, with caching).

As differences in performance were more pronounced for the larger set sizes, we tried more instances with $n = 20$, which is the largest for which we could obtain the optimal results in reasonable time. The results are shown as a cumulative average plot (Figure 5). While the value of information for all greedy algorithms is still close to the optimal value, the compound greedy algorithm again is slighly better, averaging 0.99 of the optimal net VPI, while the additive and rate greedy averaged roughly 0.98 of the optimal net VPI. The compound-greedy algorithm showed an improvement over the simple greedy algorithms whether or not condition C1 held.

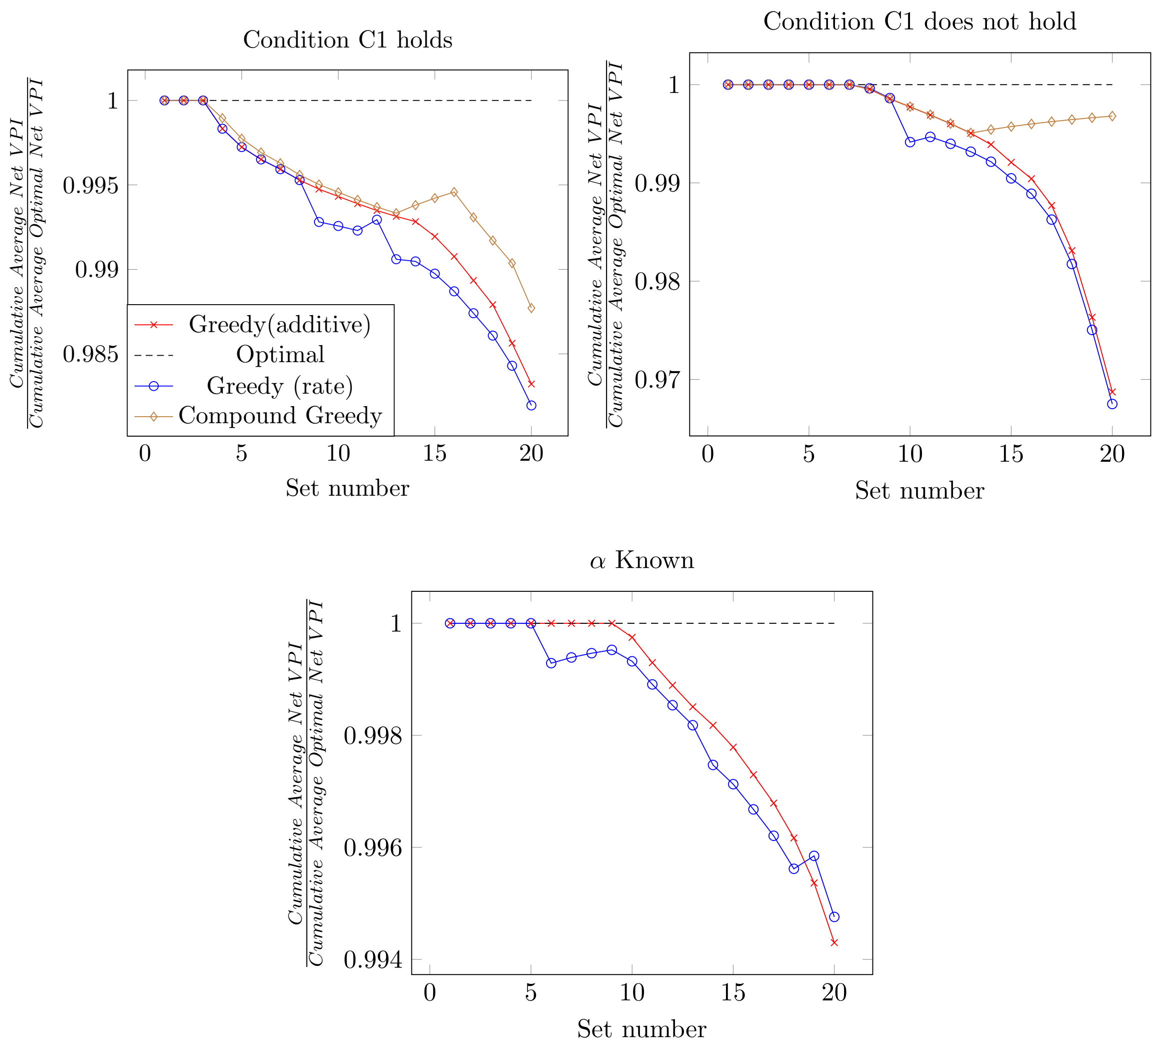

Figure 5: Comparison of net VPI for item set size 20

4. Conclusion

We have examined cases where the batch value of perfect information is submodular in the selection problem, mostly in the case where the item utility distributions are independent. We have shown that a resulting optimization problem is NP-hard, even in such restricted cases. Nevertheless, greedy optimization algorithms seem to achieve good results in practice. The theoretical results suggest that greedy algorithms should be supplemented by examining sets that include the currently best item, even if its individual VPI is zero, and this is supported by empirical evidence.

We suggest that deviations from submodularity indicate points where the greedy and myopic optimization schemes can be improved w.r.t. net VPI, at relatively little computational cost. As such, the simple method suggested in this paper complements the idea of "blinkered VOI" (Hay et al., 2012). Our motivation for this work comes from meta-reasoning in search, where the information is gathered by search actions, and solving a selection problem is a first step that suggests a way to proceed at the first level in the search tree. Generalization of these methods to selecting computations at deeper levels of the search tree is a non-trivial issue for future work.

Acknowledgments

Partially supported by ISF grant 417/13, and by the Frankel fund for Computer Science. We thank the anonymous reviewers for comments that shortened some theorem proofs, and other suggestions that increase the applicability of the results presented in this paper.

Appendix A

Proof (of Theorem 2): By reduction from the Knapsack problem (Garey & Johnson, 1979); see problem number [MP9]. We re-stated the problem below for convenience.

Definition 4 (Knapsack problem)**.** *Given a set of items $\mathcal{S} = \{s_1, ..., s_n\}$, each with a positive integer weight w_i and a positive integer value v_i, a weight limit W and a target value V, is there a subset $S \in \mathcal{S}$ such that the total weight of S is at most W, and the total value of the elements of S is at least V?*

We assume w.l.o.g. that $v_i < V$ for all items, as items that violate this solve the Knapsack problem trivially.

In reducing Knapsack to PBSP, each item in $\mathcal{I}$ (except for a special item s_α) in the selection problem will stand for the respective element in the Knapsack problem. As this is a simple one-to-one mapping, for the sake of simplicity we therefore abuse the notation and treat the non-s_α items as if they actually were the respective elements from $\mathcal{S}$ in the Knapsack problem. The distributions of values and costs are defined as follows: let $H = \max\limits_{1 \le i \le n} \{v_i\}$, and $\varepsilon = \frac{1}{2H^2n^3}$. Then let $C = W$, $U = \varepsilon(V - \frac{1}{2})$, and let the distributions and measurement costs be as follows:

- For X_α we have $u_\alpha = 0$ with probability 1. The cost C_α is irrelevant (because the exact value of u_α is already known) and can be taken to be 0.

- For every other item, we have a binary-valued distribution: $P(X_i = 1) = \varepsilon v_i$, and $P(X_i = -1) = 1 - (\varepsilon v_i)$. The measurement cost of these items is given by $C_i = w_i$.

Note that indeed the current best item is s_α, because $\varepsilon v_i < \frac{1}{2}$ for all $1 \le i \le n$. Therefore, each item s_i becomes better than s_α if and only if s_i is observed to have a positive utility. We now show that a subset $S \subseteq \mathcal{S}$ solves the Knapsack problem if and only if S solves the PBSP. Let $m = |S| \le n$, and in order to simplify the notation below, we assume w.l.o.g. that $S = \{s_1, s_2, ..., s_m\}$.

($\Rightarrow$) Let S be a solution to the Knapsack problem. We have $\sum_{i=1}^{m} w_i \le W = C$, so S satisfies the budget constraint in the PBSP. Denote the probability that at least one of the items in S has value 1 by $P(S)$. Since by construction the expected utility for a set of measurements S is exactly $P(S)$, it is sufficient to show that $P(S) \geq U$. Since the value distributions are jointly independent, we have:

$$P(S) = \sum_{i=1}^{m} P(X_i = 1) \prod_{j=1}^{i-1} (1 - P(X_j = 1)) = \sum_{i=1}^{m} \varepsilon v_i \prod_{j=1}^{i-1} (1 - \varepsilon v_j)$$

Re-arranging $P(S)$ into sums according to powers of ε, we get:

$$P(S) = \sum_{i=1}^{m} \varepsilon^i (-1)^{(i+1)} \sum_{\{N \subseteq [1..m] \wedge |N| = i\}} \prod_{k \in N} v_k \geq \varepsilon \sum_{j=1}^{m} v_j - \sum_{i'=1}^{\lfloor m/2 \rfloor} \varepsilon^{2i'} \sum_{\{N \subseteq [1..m] \wedge |N| = 2i'\}} \prod_{k \in N} v_k$$

where the inequality is due to dropping all the terms for odd i (which are positive), except for $i = 1$. Now, for each i', the number of elements in N is clearly $\binom{m}{2i'}$, which is bounded by $m^{2i'}$, and thus by $n^{2i'}$. Since by definition we also have $v_k \leq H$ for all k, we get:

$$\begin{aligned}
P(S) &\geq \varepsilon \sum_{j=1}^{m} v_j - \sum_{i'=1}^{\lfloor m/2 \rfloor} (\varepsilon H n)^{2i'} \\
&= \varepsilon \sum_{j=1}^{m} v_j - \sum_{i'=1}^{\lfloor m/2 \rfloor} (\frac{1}{2H^2 n^3} H n)^{2i'} = \varepsilon \sum_{j=1}^{m} v_j - \sum_{i'=1}^{\lfloor m/2 \rfloor} (\frac{1}{2Hn^2})^{2i'} \\
&> \varepsilon \sum_{j=1}^{m} v_j - \frac{n}{4H^2 n^4} = \varepsilon \sum_{j=1}^{m} v_j - \frac{1}{4H^2 n^3} = \varepsilon \sum_{j=1}^{m} v_j - \frac{\varepsilon}{2} \\
&= \varepsilon (V - \frac{1}{2}) = U
\end{aligned}$$

where the last equality follows from S being a solution to the Knapsack problem. Therefore, S is a solution to the PBSP.

($\Leftarrow$) Let S be a solution to the PBSP. and thus $\sum_{1=1}^{m} C_i \leq C = W$, so S obeys the weight limitation of the Knapsack problem. It is thus sufficient to show that $\sum_{i=1}^{m} v_i \geq V$. As above, we have:

$$P(S) = \sum_{i=1}^{m} \varepsilon^i (-1)^{(i+1)} \sum_{\{N \subseteq [1..m] \wedge |N| = i\}} \prod_{k \in N} v_k \leq \varepsilon \sum_{j=1}^{m} v_j + \sum_{i'=1}^{\lceil m/2 \rceil - 1} \varepsilon^{2i'+1} \sum_{\{N \subseteq [1..m] \wedge |N| = 2i'+1\}} \prod_{k \in N} v_k$$

where the inequality is due to dropping all the terms for even i (which are negative). Now, for each i', the number of elements in N is clearly $\binom{m}{2i'+1}$, which is bounded by $m^{2i'+1}$, and thus by $n^{2i'+1}$. Since by definition we also have $v_k \leq H$ for all k, we get:

$$
\begin{aligned}
P(S) &\leq \varepsilon \sum_{j=1}^{m} v_j + \sum_{i'=1}^{\lceil m/2 \rceil - 1} (\varepsilon H n)^{2i'+1} \\
&= \varepsilon \sum_{j=1}^{m} v_j + \sum_{i'=1}^{\lceil m/2 \rceil - 1} (\frac{1}{2H^2 n^3} H n)^{2i'+1} = \varepsilon \sum_{j=1}^{m} v_j + \sum_{i'=1}^{\lceil m/2 \rceil - 1} (\frac{1}{2Hn^2})^{2i'+1} \\
&< \varepsilon \sum_{j=1}^{m} v_j + \frac{n}{8H^3 n^6} = \varepsilon \sum_{j=1}^{m} v_j + \frac{1}{8H^3 n^5} \\
&< \varepsilon \sum_{j=1}^{m} v_j + \frac{1}{8H^2 n^3} = \varepsilon \sum_{j=1}^{m} v_j + \frac{\varepsilon}{4} = \varepsilon (\sum_{j=1}^{m} v_j + \frac{1}{4})
\end{aligned}
$$

Since S is a solution to PBSP, we have $P(S) \geq U = \varepsilon(V - \frac{1}{2})$, and thus $V \leq \sum_{j=1}^{m} v_j + \frac{3}{4}$. As both V and $\sum_{j=1}^{m} v_j$ are positive integers, we also have $\sum_{j=1}^{m} v_j \geq V$. Therefore, S is a solution to the Knapsack problem. □

References

Azimi, J., Fern, X., & Fern, A. (2016). Budgeted optimization with constrained experiments. *Journal of Artif. Intell. Research*, *56*, 119–152.

Chadeand, H., & Schlee, E. E. (2001). Another look at the Radner–Stiglitz nonconcavity in the value of information. *Journal of Economic Theory*, *107*, 421–452.

Cortez, P., Cerdeira, A., Almeida, F., Matos, T., & Reis, J. (2009). Modeling wine preferences by data mining from physicochemical properties. *Decision Support Systems*, *47*(4), 547–553.

Cortez, P. (2009). Wine Quality Data Set. `https://archive.ics.uci.edu/ml/datasets/Wine+Quality`. [Online; UCI Machine Learning Repository].

Frazier, P. I. (2012). Optimization via simulation with Bayesian statistics and dynamic programming. In Rose, O., & Uhrmacher, A. M. (Eds.), *Winter Simulation Conference, WSC '12, Berlin, Germany, December 9-12, 2012*, pp. 7:1–7:16. WSC.

Frazier, P. I., & Powell, W. B. (2010). Paradoxes in learning and the marginal value of information. *Decision Analysis*, *7*(4), 378–403.

Garey, M. R., & Johnson, D. S. (1979). *Computers and Intractability, A Guide to the Theory of NP-completeness*, p. 190. W. H. Freeman and Co.

Hay, N., Russell, S. J., Tolpin, D., & Shimony, S. E. (2012). Selecting computations: Theory and applications. In de Freitas, N., & Murphy, K. P. (Eds.), *Proceedings of the Twenty-Eighth Conference on Uncertainty in Artificial Intelligence, Catalina Island, CA, USA, August 14-18, 2012*, pp. 346–355. AUAI Press.

Krause, A., & Guestrin, C. (2009). Optimal value of information in graphical models. *J. Artif. Intell. Res. (JAIR)*, *35*, 557–591.

Krause, A., & Guestrin, C. (2011). Submodularity and its applications in optimized information gathering. *ACM TIST*, *2*(4), 32.

Lara, M. D., & Gilotte, L. (2007). A tight sufficient condition for Radner–Stiglitz nonconcavity in the value of information. *Journal of Economic Theory*, *131*(1), 696–708.

Nemhauser, G. L., Wolsey, L. A., & Fisher, M. L. (1978). An analysis of approximations for maximizing submodular set functions - I. *Mathematical Programming*, *14*(1), 265–294.

Papachristoudis, G., & Fisher III, J. W. (2012). Theoretical guarantees on penalized information gathering. In *Statistical Signal Processing Workshop (SSP)*, pp. 301–304.

Radner, R., & Stiglitz, J. E. (1984). A nonconcavity in the value of information. *Bayesian Models in Economic Theory*, *5*, 33–52.

Radovilsky, Y., Shattah, G., & Shimony, S. E. (2006). Efficient deterministic approximation algorithms for non-myopic value of information in graphical models. In *Proceedings of the IEEE International Conference on Systems, Man and Cybernetics, Taipei, Taiwan, October 8-11, 2006*, pp. 2559–2564. IEEE.

Radovilsky, Y., & Shimony, S. E. (2008). Observation subset selection as local compilation of performance profiles. In McAllester, D. A., & Myllymäki, P. (Eds.), *UAI 2008, Proceedings of the 24th Conference in Uncertainty in Artificial Intelligence, Helsinki, Finland, July 9-12, 2008*, pp. 460–467. AUAI Press.

Raiffa, H., & Schlaifer, R. (2000). *Applied Statistical Decision Theory*. Wiley.

Reches, S., Gal, Y. K., & Kraus, S. (2013). Efficiently gathering information in costly domains. *Decision Support Systems*, *55*(1), 326–335.

Russell, S. J., & Wefald, E. (1991a). *Do the right thing - studies in limited rationality*. MIT Press.

Russell, S. J., & Wefald, E. (1991b). Principles of metareasoning. *Artif. Intell.*, *49*(1-3), 361–395.

Silverman, B. W. (1986). *Density Estimation for Statistics and Data Analysis*. Chapman & Hall, London.

Swisher, J. R., Jacobson, S. H., & Yücesan, E. (2003). Discrete-event simulation optimization using ranking, selection, and multiple comparison procedures: A survey. *ACM Trans. Model. Comput. Simul.*, *13*(2), 134–154.

Tolpin, D., & Shimony, S. E. (2012). Semimyopic measurement selection for optimization under uncertainty. *IEEE Transactions on Systems, Man, and Cybernetics, Part B*, *42*(2), 565–579.