
Near-Optimal Bayesian Active Learning with Noisy Observations

Daniel Golovin
Caltech

Andreas Krause
Caltech

Debajyoti Ray
Caltech

Abstract

We tackle the fundamental problem of Bayesian active learning with noise, where we need to adaptively select from a number of expensive tests in order to identify an unknown hypothesis sampled from a known prior distribution. In the case of noise-free observations, a greedy algorithm called generalized binary search (GBS) is known to perform near-optimally. We show that if the observations are noisy, perhaps surprisingly, GBS can perform very poorly. We develop EC², a novel, greedy active learning algorithm and prove that it is competitive with the optimal policy, thus obtaining the first competitiveness guarantees for Bayesian active learning with noisy observations. Our bounds rely on a recently discovered diminishing returns property called adaptive submodularity, generalizing the classical notion of submodular set functions to adaptive policies. Our results hold even if the tests have non-uniform cost and their noise is correlated. We also propose EFFECX-TIVE, a particularly fast approximation of EC², and evaluate it on a Bayesian experimental design problem involving human subjects, intended to tease apart competing economic theories of how people make decisions under uncertainty.

1 Introduction

How should we perform experiments to determine the most accurate scientific theory among competing candidates, or choose among expensive medical procedures to accurately determine a patient's condition, or select which labels to obtain in order to determine the hypothesis that minimizes generalization error? In all these applications, we have to sequentially select among a set of noisy, expensive observations (outcomes of experiments, medical tests, expert labels) in order to determine which hypothesis (theory, diagnosis, classifier) is most accurate. This fundamental problem has been studied in a number of areas, including statistics [17], decision theory [13], machine learning [19, 7] and others. One way to formalize such active learning problems is *Bayesian experimental design* [6], where one assumes a prior on the hypotheses, as well as probabilistic assumptions on the outcomes of tests. The goal then is to determine the correct hypothesis while minimizing the cost of the experimentation. Unfortunately, finding this optimal policy is not just NP-hard, but also hard to approximate [5]. Several heuristic approaches have been proposed that perform well in some applications, but do not carry theoretical guarantees (e.g., [18]). In the case where observations are *noise-free*¹, a simple algorithm, *generalized binary search*²(GBS) run on a modified prior, is guaranteed to be competitive with the optimal policy; the expected number of queries is a factor of $O(\log n)$ (where n is the number of hypotheses) more than that of the optimal policy [15], which matches lower bounds up to constant factors [5].

The important case of *noisy* observations, however, as present in most applications, is much less well understood. While there are some recent positive results in understanding the label complexity of noisy active learning [19, 1], to our knowledge, so far there are no algorithms that are provably competitive with the optimal sequential policy, except in very restricted settings [16]. In this paper, we

¹This case is known as the *Optimal Decision Tree* (ODT) problem.

²GBS greedily selects tests to maximize, in expectation over the test outcomes, the prior probability mass of eliminated hypotheses (i.e., those with zero posterior probability, computed w.r.t. the observed test outcomes).

introduce a general formulation of Bayesian active learning with noisy observations that we call the *Equivalence Class Determination* problem. We show that, perhaps surprisingly, generalized binary search performs poorly in this setting, as do greedily (myopically) maximizing the information gain (measured w.r.t. the distribution on equivalence classes) or the decision-theoretic value of information. This motivates us to introduce a novel active learning criterion, and use it to develop a greedy active learning algorithm called the **E**quivalence **C**lass **E**dge **C**utting algorithm (EC^2), whose expected cost is competitive to that of the optimal policy. Our key insight is that our new objective function satisfies *adaptive submodularity* [9], a natural diminishing returns property that generalizes the classical notion of submodularity to adaptive policies. Our results also allow us to relax the common assumption that the outcomes of the tests are conditionally independent given the true hypothesis. We also develop the **E**fficient **E**dge **C**utting **a**ppro**X**imate **o**bjective algorithm (EFFECXTIVE), an efficient approximation to EC^2 , and evaluate it on a Bayesian experimental design problem intended to tease apart competing theories on how people make decisions under uncertainty, including Expected Value [22], Prospect Theory [14], Mean-Variance-Skewness [12] and Constant Relative Risk Aversion [20]. In our experiments, EFFECXTIVE typically outperforms existing experimental design criteria such as information gain, uncertainty sampling, GBS, and decision-theoretic value of information. Our results from human subject experiments further reveal that EFFECXTIVE can be used as a real-time tool to classify people according to the economic theory that best describes their behaviour in financial decision-making, and reveal some interesting heterogeneity in the population.

2 Bayesian Active Learning in the Noiseless Case

In the Bayesian active learning problem, we would like to distinguish among a given set of hypotheses $\mathcal{H} = \{h_1, \dots, h_n\}$ by performing tests from a set $\mathcal{T} = \{1, \dots, N\}$ of possible tests. Running test t incurs a cost of $c(t)$ and produces an outcome from a finite set of outcomes $\mathcal{X} = \{1, 2, \dots, \ell\}$. We let H denote the random variable which equals the true hypothesis, and model the outcome of each test t by a random variable X_t taking values in \mathcal{X} . We denote the observed outcome of test t by x_t . We further suppose we have a prior distribution P modeling our assumptions on the joint probability $P(H, X_1, \dots, X_N)$ over the hypotheses and test outcomes. In the noiseless case, we assume that the outcome of each test is deterministic given the true hypothesis, i.e., for each $h \in \mathcal{H}$, $P(X_1, \dots, X_N \mid H = h)$ is a deterministic distribution. Thus, each hypothesis h is associated with a particular vector of test outcomes. We assume, w.l.o.g., that no two hypotheses lead to the same outcomes for all tests. Thus, if we perform all tests, we can uniquely determine the true hypothesis. However in most applications we will wish to avoid performing every possible test, as this is prohibitively expensive. Our goal is to find an adaptive policy for running tests that allows us to determine the value of H while minimizing the cost of the tests performed. Formally, a policy π (also called a conditional plan) is a partial mapping π from partial observation vectors $\mathbf{x}_{\mathcal{A}}$ to tests, specifying which test to run next (or whether we should stop testing) for any observation vector $\mathbf{x}_{\mathcal{A}}$. Hereby, $\mathbf{x}_{\mathcal{A}} \in \mathcal{X}^{\mathcal{A}}$ is a vector of outcomes indexed by a set of tests $\mathcal{A} \subseteq \mathcal{T}$ that we have performed so far³ (e.g., the set of labeled examples in active learning, or outcomes of a set of medical tests that we ran). After having made observations $\mathbf{x}_{\mathcal{A}}$, we can rule out inconsistent hypotheses. We denote the set of hypotheses consistent with event Λ (often called the *version space* associated with Λ) by $\mathcal{V}(\Lambda) := \{h \in \mathcal{H} : P(h \mid \Lambda) > 0\}$. We call a policy *feasible* if it is guaranteed to uniquely determine the correct hypothesis. That is, upon termination with observation $\mathbf{x}_{\mathcal{A}}$, it must hold that $|\mathcal{V}(\mathbf{x}_{\mathcal{A}})| = 1$. We can define the expected cost of a policy π by

$$c(\pi) := \sum_h P(h) c(\mathcal{T}(\pi, h))$$

where $\mathcal{T}(\pi, h) \subseteq \mathcal{T}$ is the set of tests run by policy π in case $H = h$. Our goal is to find a feasible policy π^* of minimum expected cost, i.e.,

$$\pi^* = \arg \min \{c(\pi) : \pi \text{ is feasible}\} \quad (2.1)$$

A policy π can be naturally represented as a decision tree T^π , and thus problem (2.1) is often called the *Optimal Decision Tree* (ODT) problem.

Unfortunately, obtaining an approximate policy π for which $c(\pi) \leq c(\pi^*) \cdot o(\log(n))$ is NP-hard [5]. Hence, various heuristics are employed to solve the Optimal Decision Tree problem and its variants. Two of the most popular heuristics are to select tests greedily to maximize the *information gain* (IG)

³Formally we also require that $(x_t)_{t \in \mathcal{B}} \in \text{dom}(\pi)$ and $\mathcal{A} \subseteq \mathcal{B}$, implies $(x_t)_{t \in \mathcal{A}} \in \text{dom}(\pi)$ (c.f., [9]).

conditioned on previous test outcomes, and *generalized binary search* (GBS). Both heuristics are greedy, and after having made observations \mathbf{x}_A will select

$$t^* = \arg \max_{t \in \mathcal{T}} \Delta_{\text{Alg}}(t | \mathbf{x}_A) / c(t),$$

where $\text{Alg} \in \{\text{IG}, \text{GBS}\}$. Here, $\Delta_{\text{IG}}(t | \mathbf{x}_A) := \mathbb{H}(\mathbf{X}_{\mathcal{T}} | \mathbf{x}_A) - \mathbb{E}_{x_t \sim X_t | \mathbf{x}_A} [\mathbb{H}(\mathbf{X}_{\mathcal{T}} | \mathbf{x}_A, x_t)]$ is the marginal information gain measured with respect to the Shannon entropy $\mathbb{H}(\mathbf{X}) := \mathbb{E}_{\mathbf{x}}[-\log_2 P(\mathbf{x})]$, and $\Delta_{\text{GBS}}(t | \mathbf{x}_A) := P(\mathcal{V}(\mathbf{x}_A)) - \sum_{x \in \mathcal{X}} P(X_t = x | \mathbf{x}_A) P(\mathcal{V}(\mathbf{x}_A, X_t = x))$ is the expected reduction in version space probability mass. Thus, both heuristics greedily chooses the test that maximizes the benefit-cost ratio, measured with respect to their particular benefit functions. They stop after running a set of tests \mathcal{A} such that $|\mathcal{V}(\mathbf{x}_A)| = 1$, i.e., once the true hypothesis has been uniquely determined.

It turns out that for the (noiseless) Optimal Decision Tree problem, these two heuristics are equivalent [23], as can be proved using the chain rule of entropy. Interestingly, despite its myopic nature GBS has been shown [15, 7, 11, 9] to obtain near-optimal expected cost: the strongest known bound is $c(\pi_{\text{GBS}}) \leq c(\pi^*) (\ln(1/p_{\min}) + 1)$ where $p_{\min} := \min_{h \in \mathcal{H}} P(h)$. Let $\mathbf{x}_S(h)$ be the unique vector $\mathbf{x}_S \in \mathcal{X}^S$ such that $P(\mathbf{x}_S | h) = 1$. The result above is proved by exploiting the fact that $f_{\text{GBS}}(S, h) := 1 - P(\mathcal{V}(\mathbf{x}_S(h))) + P(h)$ is *adaptive submodular* and *strongly adaptively monotone* [9]. Call \mathbf{x}_A a *subvector* of \mathbf{x}_B if $\mathcal{A} \subseteq \mathcal{B}$ and $P(\mathbf{x}_B | \mathbf{x}_A) > 0$. In this case we write $\mathbf{x}_A \prec \mathbf{x}_B$. A function $f : 2^{\mathcal{T}} \times \mathcal{H}$ is called *adaptive submodular* w.r.t. a distribution P , if for any $\mathbf{x}_A \prec \mathbf{x}_B$ and any test t it holds that $\Delta(t | \mathbf{x}_A) \geq \Delta(t | \mathbf{x}_B)$, where

$$\Delta(t | \mathbf{x}_A) := \mathbb{E}_H[f(\mathcal{A} \cup \{t\}, H) - f(\mathcal{A}, H) | \mathbf{x}_A].$$

Thus, f is *adaptive submodular* if the expected marginal benefits $\Delta(t | \mathbf{x}_A)$ of adding a new test t can only decrease as we gather more observations. f is called *strongly adaptively monotone* w.r.t. P if, informally, “observations never hurt” with respect to the expected reward. Formally, for all \mathcal{A} , all $t \notin \mathcal{A}$, and all $x \in \mathcal{X}$ we require $\mathbb{E}_H[f(\mathcal{A}, H) | \mathbf{x}_A] \leq \mathbb{E}_H[f(\mathcal{A} \cup \{t\}, H) | \mathbf{x}_A, X_t = x]$.

The performance guarantee for GBS follows from the following general result about the greedy algorithm for adaptive submodular functions (applied with $Q = 1$ and $\eta = p_{\min}$):

Theorem 1 (Theorem 10 of [9] with $\alpha = 1$). *Suppose $f : 2^{\mathcal{T}} \times \mathcal{H} \rightarrow \mathbb{R}_{\geq 0}$ is adaptive submodular and strongly adaptively monotone with respect to P and there exists Q such that $f(\mathcal{T}, h) = Q$ for all h . Let η be any value such that $f(S, h) > Q - \eta$ implies $f(S, h) = Q$ for all sets S and hypotheses h . Then for self-certifying instances the adaptive greedy policy π satisfies $c(\pi) \leq c(\pi^*) \left(\ln \left(\frac{Q}{\eta} \right) + 1 \right)$.*

The technical requirement that instances be *self-certifying* means that the policy will have proof that it has obtained the maximum possible objective value, Q , immediately upon doing so. It is not difficult to show that this is the case with the instances we consider in this paper. We refer the interested reader to [9] for more detail.

In the following sections, we will use the concept of adaptive submodularity to provide the first approximation guarantees for Bayesian active learning with noisy observations.

3 The Equivalence Class Determination Problem and the EC² Algorithm

We now wish to consider the Bayesian active learning problem where tests can have noisy outcomes. Our general strategy is to reduce the problem of noisy observations to the noiseless setting. To gain intuition, consider a simple model where tests have binary outcomes, and we know that the outcome of exactly one test, chosen uniformly at random unbeknown to us, is flipped. If any pair of hypotheses $h \neq h'$ differs by the outcome of at least three tests, we can still uniquely determine the correct hypothesis after running all tests. In this case we can reduce the noisy active learning problem to the noiseless setting by, for each hypothesis, creating N “noisy” copies, each obtained by flipping the outcome of one of the N tests. The modified prior P' would then assign mass $P'(h') = P(h)/N$ to each noisy copy h' of h . The conditional distribution $P'(\mathbf{X}_{\mathcal{T}} | h')$ is still deterministic (obtained by flipping the outcome of one of the tests). Thus, each hypothesis h_i in the original problem is now associated with a set \mathcal{H}_i of hypotheses in the modified problem instance. However, instead of selecting tests to determine which noisy copy has been realized, we only care which set \mathcal{H}_i is realized.

The Equivalence Class Determination problem (ECD). More generally, we introduce the *Equivalence Class Determination problem*⁴, where our set of hypotheses \mathcal{H} is partitioned into a set of m equivalence classes $\{\mathcal{H}_1, \dots, \mathcal{H}_m\}$ so that $\mathcal{H} = \bigsqcup_{i=1}^m \mathcal{H}_i$, and the goal is to determine which class \mathcal{H}_i the true hypothesis lies in. Formally, upon termination with observations \mathbf{x}_A we require that $\mathcal{V}(\mathbf{x}_A) \subseteq \mathcal{H}_i$ for some i . As with the ODT problem, the goal is to minimize the expected cost of the tests, where the expectation is taken over the true hypothesis sampled from P . In §4, we will show how the Equivalence Class Determination problem arises naturally from Bayesian experimental design problems in probabilistic models.

Given the fact that GBS performs near-optimally on the Optimal Decision Tree problem, a natural approach to solving ECD would be to run GBS until the termination condition is met. Unfortunately, and perhaps surprisingly, GBS can perform very poorly on the ECD problem. Consider an instance with a uniform prior over n hypotheses, h_1, \dots, h_n , and two equivalence classes $\mathcal{H}_1 := \{h_i : 1 \leq i < n\}$ and $\mathcal{H}_2 := \{h_n\}$. There are tests $\mathcal{T} = \{1, \dots, n\}$ such that $h_i(t) = \mathbf{1}[i = t]$, all of unit cost. Hereby, $\mathbf{1}[\Lambda]$ is the indicator variable for event Λ . In this case, the optimal policy only needs to select test n , however GBS may select tests $1, 2, \dots, n$ in order until running test t , where $H = h_t$ is the true hypothesis. Given our uniform prior, it takes $n/2$ tests in expectation until this happens, so that GBS pays, in expectation, $n/2$ times the optimal expected cost in this instance.

The poor performance of GBS in this instance may be attributed to its lack of consideration for the equivalence classes. Another natural heuristic would be to run the greedy information gain policy, only with the entropy measured with respect to the probability distribution on *equivalence classes* rather than hypotheses. Call this policy π_{IG} . It is clearly aware of the equivalence classes, as it adaptively and myopically selects tests to reduce the uncertainty of the realized class, measured w.r.t. the Shannon entropy. However, we can show there are instances in which it pays $\Omega(n/\log(n))$ times the optimal cost, even under a uniform prior. See the long version of this paper [10] for details.

The EC² algorithm. The reason why GBS fails is because reducing the version space mass does not necessarily facilitate differentiation among the classes \mathcal{H}_i . The reason why π_{IG} fails is that there are complementarities among tests; a set of tests can be far better than the sum of its parts. Thus, we would like to optimize an objective function that encourages differentiation among classes, but lacks complementarities. We adopt a very elegant idea from Dasgupta [8], and define weighted edges between hypotheses that we aim to distinguish between. However, instead of introducing edges between arbitrary pairs of hypotheses (as done in [8]), we only introduce edges between hypotheses in different classes. Tests will allow us to cut edges inconsistent with their outcomes, and we aim to eliminate all inconsistent edges while minimizing the expected cost incurred. We now formalize this intuition.

Specifically, we define a set of edges $\mathcal{E} = \cup_{1 \leq i < j \leq m} \{\{h, h'\} : h \in \mathcal{H}_i, h' \in \mathcal{H}_j\}$, consisting of all (unordered) pairs of hypotheses belonging to distinct classes. These are the edges that must be *cut*, by which we mean for any edge $\{h, h'\} \in \mathcal{E}$, at least one hypothesis in $\{h, h'\}$ must be ruled out (i.e., eliminated from the version space). Hence, a test t run under true hypothesis h is said to cut edges $\mathcal{E}_t(h) := \{\{h', h''\} : h'(t) \neq h(t) \text{ or } h''(t) \neq h(t)\}$. See Fig. 1(a) for an illustration. We define a weight function $w : \mathcal{E} \rightarrow \mathbb{R}_{\geq 0}$ by $w(\{h, h'\}) := P(h) \cdot P(h')$. We extend the weight function to an additive (modular) function on sets of edges in the natural manner, i.e., $w(\mathcal{E}') := \sum_{e \in \mathcal{E}'} w(e)$. The objective f_{EC} that we will greedily maximize is then defined as the weight of the edges cut (EC):

$$f_{EC}(\mathcal{A}, h) := w\left(\bigcup_{t \in \mathcal{A}} \mathcal{E}_t(h)\right) \quad (3.1)$$

The key insight that allows us to prove approximation guarantees for f_{EC} is that f_{EC} shares the same beneficial properties that make f_{GBS} amenable to efficient greedy optimization. The proof of this fact, as stated in Proposition 2, can be found in the long version of this paper [10].

Proposition 2. *The objective f_{EC} is strongly adaptively monotone and adaptively submodular.*

Based on the objective f_{EC} , we can calculate the marginal benefits for test t upon observations \mathbf{x}_A as

$$\Delta_{EC}(t | \mathbf{x}_A) := \mathbb{E}_H[f_{EC}(\mathcal{A} \cup \{t\}, H) - f_{EC}(\mathcal{A}, H) \mid \mathbf{x}_A].$$

We call the adaptive policy π_{EC} that, after observing \mathbf{x}_A , greedily selects test $t^* \in \arg \max_t \Delta_{EC}(t | \mathbf{x}_A) / c(t)$, the EC² algorithm (for *equivalence class edge cutting*).

⁴Bellala et al. simultaneously studied ECD [2], and, like us, used it to model active learning with noise [3]. They developed an extension of GBS for ECD. We defer a detailed comparison of our approaches to future work.




Figure 1: (a) An instance of Equivalence Class Determination with binary test outcomes, shown with the set of edges that must be cut, and depicting the effects of test i under different outcomes. (b) The graphical model underlying our error model.

Note that these instances are self-certifying, because we obtain maximum objective value if and only if the version space lies within an equivalence class, and the policy can certify this condition when it holds. So we can apply Theorem 1 to show EC^2 obtains a $\ln(Q/\eta) + 1$ approximation to Equivalence Class Determination. Hereby, $Q = w(\mathcal{E}) = 1 - \sum_i (P(h \in \mathcal{H}_i))^2 \leq 1$ is the total weight of all edges that need to be cut, and $\eta = \min_{e \in \mathcal{E}} w(e) \geq p_{\min}^2$ is a bound on the minimum weight among all edges. We have the following result:

Theorem 3. *Suppose $P(h)$ is rational for all $h \in \mathcal{H}$. For the adaptive greedy policy π_{EC} implemented by EC^2 it holds that*

$$c(\pi_{\text{EC}}) \leq (2 \ln(1/p_{\min}) + 1)c(\pi^*),$$

where $p_{\min} := \min_{h \in \mathcal{H}} P(h)$ is the minimum prior probability of any hypothesis, and π^* is the optimal policy for the Equivalence Class Determination problem.

In the case of unit cost tests, we can apply a technique of Kosaraju et al. [15], originally developed for the GBS algorithm, to improve the approximation guarantee to $\mathcal{O}(\log n)$ by applying EC^2 with a modified prior distribution. We defer details to the full version of this paper.

4 Bayesian Active Learning with Noise and the EFFECXTIVE Algorithm

We now address the case of noisy observations, using ideas from §3. With noisy observations, the conditional distribution $P(X_1, \dots, X_N | h)$ is no longer deterministic. We model the noise using an additional random variable Θ . Fig. 1(b) depicts the underlying graphical model. The vector of test outcomes $\mathbf{x}_{\mathcal{T}}$ is assumed to be an arbitrary, deterministic function $\mathbf{x}_{\mathcal{T}} : \mathcal{H} \times \text{supp}(\Theta) \rightarrow \mathcal{X}^N$; hence $\mathbf{X}_{\mathcal{T}} | h$ is distributed as $\mathbf{x}_{\mathcal{T}}(h, \Theta_h)$ where Θ_h is distributed as $P(\theta | h)$. For example, there might be up to $s = |\text{supp}(\Theta)|$ ways any particular disease could manifest itself, with different patients with the same disease suffering from different symptoms.

In cases where it is always possible to identify the true hypothesis, i.e., $\mathbf{x}_{\mathcal{T}}(h, \theta) \neq \mathbf{x}_{\mathcal{T}}(h', \theta')$ for all $h \neq h'$ and all $\theta, \theta' \in \text{supp}(\Theta)$, we can reduce the problem to Equivalence Class Determination with hypotheses $\{\mathbf{x}_{\mathcal{T}}(h, \theta) : h \in \mathcal{H}, \theta \in \text{supp}(\Theta)\}$ and equivalence classes $\mathcal{H}_i := \{\mathbf{x}_{\mathcal{T}}(h_i, \theta) : \theta \in \text{supp}(\Theta)\}$ for all i . Then Theorem 3 immediately yields that the approximation factor of EC^2 is at most $2 \ln(1/\min_{h, \theta} P(h, \theta)) + 1$, where the minimum is taken over all (h, θ) in the support of P . In the unit cost case, running EC^2 with a modified prior à la Kosaraju et al. [15] allows us to obtain an $\mathcal{O}(\log |\mathcal{H}| + \log |\text{supp}(\Theta)|)$ approximation factor. Note this model allows us to incorporate noise with complex correlations.

However, a major challenge when dealing with noisy observations is that it is not always possible to distinguish distinct hypotheses. Even after we have run all tests, there will generally still be uncertainty about the true hypothesis, i.e., the posterior distribution $P(H | \mathbf{x}_{\mathcal{T}})$ obtained using Bayes' rule may still assign non-zero probability to more than one hypothesis. If so, uniquely determining the true hypothesis is not possible. Instead, we imagine that there is a set \mathcal{D} of possible decisions we may make after (adaptively) selecting a set of tests to perform and we must choose one (e.g., we must decide how to treat the medical patient, which scientific theory to adopt, or which classifier to use, given our observations). Thus our goal is to gather data to make effective decisions [13]. Formally, for any decision $d \in \mathcal{D}$ we take, and each realized hypothesis h , we incur some loss $\ell(d, h)$. Decision theory recommends, after observing $\mathbf{x}_{\mathcal{A}}$, to choose the decision d^* that minimizes the risk, i.e., the expected loss, namely $d^* \in \arg \min_d \mathbb{E}_H[\ell(d, H) | \mathbf{x}_{\mathcal{A}}]$.

A natural goal in Bayesian active learning is thus to adaptively pick observations, until we are guaranteed to make the same decision (and thus incur the same expected loss) that we would have made had we run *all* tests. Thus, we can reduce the noisy Bayesian active learning problem to the ECD problem by defining the equivalence classes over all test outcomes that lead to the same minimum risk decision. Hence, for each decision $d \in \mathcal{D}$, we define

$$\mathcal{H}_d := \{\mathbf{x}_{\mathcal{T}} : d = \arg \min_{d'} \mathbb{E}_H[\ell(d', H) \mid \mathbf{x}_{\mathcal{T}}]\}. \quad (4.1)$$

If multiple decisions minimize the risk for a particular $\mathbf{x}_{\mathcal{T}}$, we break ties arbitrarily. Identifying the best decision $d \in \mathcal{D}$ then amounts to identifying which equivalence class \mathcal{H}_d contains the realized vector of outcomes, which is an instance of ECD.

One common approach to this problem is to myopically pick tests maximizing the decision-theoretic *value of information* (VoI): $\Delta_{\text{VoI}}(t \mid \mathbf{x}_{\mathcal{A}}) := \min_d \mathbb{E}_H[\ell(d, H) \mid \mathbf{x}_{\mathcal{A}}] - \mathbb{E}_{x_t \sim X_t \mid \mathbf{x}_{\mathcal{A}}}[\min_d \mathbb{E}_H[\ell(d, H) \mid \mathbf{x}_{\mathcal{A}}, x_t]]$. The VoI of a test t is the expected reduction in the expected loss of the best decision due to the observation of x_t . However, we can show there are instances in which such a policy pays $\Omega(n/\log(n))$ times the optimal cost, even under a uniform prior on (h, θ) and with $|\text{supp}(\Theta)| = 2$. See the long version of this paper [10] for details. In contrast, on such instances EC² obtains an $\mathcal{O}(\log n)$ approximation. More generally, we have the following result for EC² as an immediate consequence of Theorem 3.

Theorem 4. *Fix hypotheses \mathcal{H} , tests \mathcal{T} with costs $c(t)$ and outcomes in \mathcal{X} , decision set \mathcal{D} , and loss function ℓ . Fix a prior $P(H, \Theta)$ and a function $\mathbf{x}_{\mathcal{T}} : \mathcal{H} \times \text{supp}(\Theta) \rightarrow \mathcal{X}^N$ which define the probabilistic noise model. Let $c(\pi)$ denote the expected cost of π incurs to find the best decision, i.e., to identify which equivalence class \mathcal{H}_d the outcome vector $\mathbf{x}_{\mathcal{T}}$ belongs to. Let π^* denote the policy minimizing $c(\cdot)$, and let π_{EC} denote the adaptive policy implemented by EC². Then it holds that*

$$c(\pi_{EC}) \leq \left(2 \ln \left(\frac{1}{p'_{\min}} \right) + 1 \right) c(\pi^*),$$

where $p'_{\min} := \min_{h \in \mathcal{H}} \{P(h, \theta) : P(h, \theta) > 0\}$.

If all tests have unit cost, by using a modified prior [15] the approximation factor can be improved to $\mathcal{O}(\log |\mathcal{H}| + \log |\text{supp}(\Theta)|)$ as in the case of Theorem 3.

The EFFECXTIVE algorithm. For some noise models, Θ may have exponentially-large support. In this case reducing Bayesian active learning with noise to Equivalence Class Determination results in instances with exponentially-large equivalence classes. This makes running EC² on them challenging, since explicitly keeping track of the equivalence classes is impractical. To overcome this challenge, we develop EFFECXTIVE, a particularly efficient algorithm which approximates EC².

For clarity, we only consider the 0–1 loss, i.e., our goal is to find the most likely hypothesis (MAP estimate) given all the data $\mathbf{x}_{\mathcal{T}}$, namely $h^*(\mathbf{x}_{\mathcal{T}}) := \arg \max_h P(h \mid \mathbf{x}_{\mathcal{T}})$. Recall definition (4.1), and consider the weight of edges between distinct equivalence classes \mathcal{H}_i and \mathcal{H}_j :

$$w(\mathcal{H}_i \times \mathcal{H}_j) = \sum_{\mathbf{x}_{\mathcal{T}} \in \mathcal{H}_i, \mathbf{x}'_{\mathcal{T}} \in \mathcal{H}_j} P(\mathbf{x}_{\mathcal{T}})P(\mathbf{x}'_{\mathcal{T}}) = \left(\sum_{\mathbf{x}_{\mathcal{T}} \in \mathcal{H}_i} P(\mathbf{x}_{\mathcal{T}}) \right) \left(\sum_{\mathbf{x}'_{\mathcal{T}} \in \mathcal{H}_j} P(\mathbf{x}'_{\mathcal{T}}) \right) = P(\mathbf{X}_{\mathcal{T}} \in \mathcal{H}_i)P(\mathbf{X}_{\mathcal{T}} \in \mathcal{H}_j).$$

In general, $P(\mathbf{X}_{\mathcal{T}} \in \mathcal{H}_i)$ can be estimated to arbitrary accuracy using a rejection sampling approach with bounded sample complexity. We defer details to the full version of the paper. Here, we focus on the case where, upon observing all tests, the hypothesis is uniquely determined, i.e., $P(H \mid \mathbf{x}_{\mathcal{T}})$ is deterministic for all $\mathbf{x}_{\mathcal{T}}$ in the support of P . In this case, it holds that $P(\mathbf{X}_{\mathcal{T}} \in \mathcal{H}_i) = P(H = h_i)$. Thus, the total weight is

$$\sum_{i \neq j} w(\mathcal{H}_i \times \mathcal{H}_j) = \left(\sum_i P(h_i) \right)^2 - \sum_i P(h_i)^2 = 1 - \sum_i P(h_i)^2.$$

This insight motivates us to use the objective function

$$\Delta_{\text{Eff}}(t \mid \mathbf{x}_{\mathcal{A}}) := \left[\sum_x P(X_t = x \mid \mathbf{x}_{\mathcal{A}}) \left(\sum_i P(h_i \mid \mathbf{x}_{\mathcal{A}}, X_t = x)^2 \right) \right] - \sum_i P(h_i \mid \mathbf{x}_{\mathcal{A}})^2,$$

which is the expected reduction in weight from the prior to the posterior distribution. Note that the weight of a distribution $1 - \sum_i P(h_i)^2$ is a monotonically increasing function of the Rényi

entropy (of order 2), which is $-\frac{1}{2} \log \sum_i P(h_i)^2$. Thus the objective Δ_{Eff} can be interpreted as a (non-standard) information gain in terms of the (exponentiated) Rényi entropy. In our experiments, we show that this criterion performs well in comparison to existing experimental design criteria, including the classical Shannon information gain. Computing $\Delta_{\text{Eff}}(t | \mathbf{x}_A)$ requires us to perform one inference task for each outcome x of X_t , and $\mathcal{O}(n)$ computations to calculate the weight for each outcome. We call the algorithm that greedily optimizes Δ_{Eff} the EFFECXTIVE algorithm (since it uses an Efficient Edge Cutting approxIMATE objective), and present pseudocode in Algorithm 1.

Input: Set of hypotheses \mathcal{H} ; Set of tests \mathcal{T} ; prior distribution P ; function f .
begin
 $\mathcal{A} \leftarrow \emptyset$;
while $\exists h \neq h' : P(h | \mathbf{x}_A) > 0$ and $P(h' | \mathbf{x}_A) > 0$ **do**
 foreach $t \in \mathcal{T} \setminus \mathcal{A}$ **do**
 $\Delta_{\text{Eff}}(t | \mathbf{x}_A) := \left[\sum_x P(X_t = x | \mathbf{x}_A) (\sum_i P(h_i | \mathbf{x}_A, X_t = x)^2) \right] - \sum_i P(h_i | \mathbf{x}_A)^2$;
 Select $t^* \in \arg \max_t \Delta_{\text{Eff}}(t | \mathbf{x}_A) / c(t)$; Set $\mathcal{A} \leftarrow \mathcal{A} \cup \{t^*\}$ and observe outcome x_{t^*} ;
 end
end

Algorithm 1: The EFFECXTIVE algorithm using the Efficient Edge Cutting approxIMATE objective.

5 Experiments

Several economic theories make claims to explain how people make decisions when the payoffs are uncertain. Here we use human subject experiments to compare four key theories proposed in literature. The uncertainty of the payoff in a given situation is represented by a *lottery* L , which is simply a random variable with a range of *payoffs* $\mathcal{L} := \{\ell_1, \dots, \ell_k\}$. For our purposes, a payoff is an integer denoting how many dollars you receive (or lose, if the payoff is negative). Fix lottery L , and let $p_i := \mathbb{P}[L = \ell_i]$. The four theories posit distinct utility functions, with agents preferring larger utility lotteries. Three of the theories have associated parameters. The *Expected Value theory* [22] posits simply $U_{EV}(L) = \mathbb{E}[L]$, and has no parameters. *Prospect theory* [14] posits $U_{PT}(L) = \sum_i f(\ell_i)w(p_i)$ for nonlinear functions $f(\ell_i) = \ell_i^\rho$, if $\ell_i \geq 0$ and $f(\ell_i) = -\lambda(-\ell_i)^\rho$, if $\ell_i < 0$, and $w(p_i) = e^{-(\log(1/p_i))^\alpha}$ [21]. The parameters $\Theta_{PT} = \{\rho, \lambda, \alpha\}$ represent risk aversion, loss aversion and probability weighing factor respectively. For portfolio optimization problems, financial economists have used value functions that give weights to different moments of the lottery [12]: $U_{MVS}(L) = w_\mu \mu - w_\sigma \sigma + w_\nu \nu$, where $\Theta_{MVS} = \{w_\mu, w_\sigma, w_\nu\}$ are the weights for the mean, standard deviation and standardized skewness of the lottery respectively. In *Constant Relative Risk Aversion theory* [20], there is a parameter $\Theta_{CRRRA} = a$ representing the level of risk aversion, and the utility posited is $U_{CRRRA}(L) = \sum_i p_i \ell_i^{1-a} / (1-a)$ if $a \neq 1$, and $U_{CRRRA}(L) = \sum_i p_i \log(\ell_i)$, if $a = 1$.

The goal is to adaptively select a sequence of tests to present to a human subject in order to distinguish which of the four theories best explains the subject’s responses. Here a test t is a pair of lotteries, (L_1^t, L_2^t) . Based on the theory that represents behaviour, one of the lotteries would be preferred to the other, denoted by a binary response $x_t \in \{1, 2\}$. The possible payoffs were fixed to $\mathcal{L} = \{-10, 0, 10\}$ (in dollars), and the distribution (p_1, p_2, p_3) over the payoffs was varied, where $p_i \in \{0.01, 0.99\} \cup \{0.1, 0.2, \dots, 0.9\}$. By considering all non-identical pairs of such lotteries, we obtained the set of possible tests.

We compare six algorithms: EFFECXTIVE, greedily maximizing Information Gain (IG), Value of Information (VOI), Uncertainty Sampling⁵ (US), Generalized Binary Search (GBS), and tests selected at Random. We evaluated the ability of the algorithms to recover the true model based on simulated responses. We chose parameter values for the theories such that they made distinct predictions and were consistent with the values proposed in literature [14]. We drew 1000 samples of the true model and fixed the parameters of the model to some canonical values, $\Theta_{PT} = \{0.9, 2.2, 0.9\}$, $\Theta_{MVS} = \{0.8, 0.25, 0.25\}$, $\Theta_{CRRRA} = 1$. Responses were generated using a softmax function, with the probability of response $x_t = 1$ given by $P(x_t = 1) = 1 / (1 + e^{U(L_2^t) - U(L_1^t)})$. Fig. 2(a) shows the performance of the 6 methods, in terms of the accuracy of recovering the true model with the number of tests. We find that US, GBS and VOI perform significantly worse than Random in the presence of noise. EFFECXTIVE outperforms InfoGain significantly, which outperforms Random.

⁵Uncertainty sampling greedily selects the test whose outcome distribution has maximum Shannon entropy.

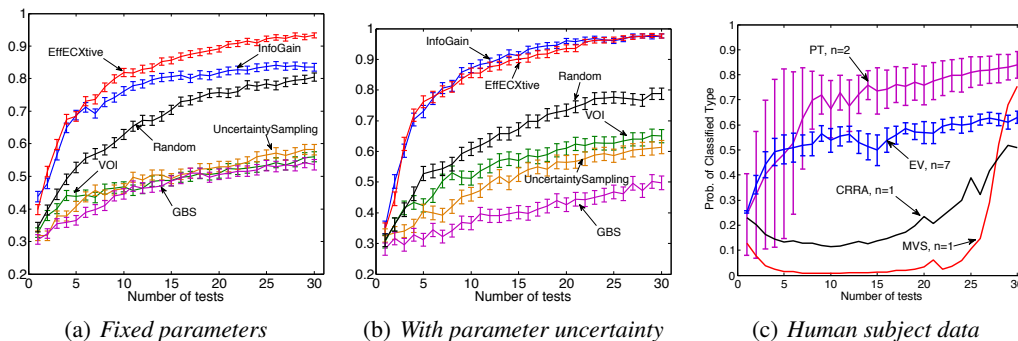


Figure 2: (a) Accuracy of identifying the true model with fixed parameters, (b) Accuracy using a grid of parameters, incorporating uncertainty in their values, (c) Experimental results: 11 subjects were classified into the theories that described their behavior best. We plot probability of classified type.

We also considered uncertainty in the values of the parameters, by setting ρ from 0.85-0.95, λ from 2.1-2.3, α from 0.9-1; w_μ from 0.8-1.0, w_σ from 0.2-0.3, w_ν from 0.2-0.3; and a from 0.9-1.0, all with 3 values per parameter. We generated 500 random samples by first randomly sampling a model and then randomly sampling parameter values. EFFECXTIVE and InfoGain outperformed Random significantly, Fig. 2(b), although InfoGain did marginally better among the two. The increased parameter range potentially poses model identifiability issues, and violates some of the assumptions behind EFFECXTIVE, decreasing its performance to the level of InfoGain.

After obtaining informed consent according to a protocol approved by the Institutional Review Board of Caltech, we tested 11 human subjects to determine which model fit their behaviour best. Laboratory experiments have been used previously to distinguish economic theories, [4], and here we used a real-time, dynamically optimized experiment that required fewer tests. Subjects were presented 30 tests using EFFECXTIVE. To incentivise the subjects, one of these tests was picked at random, and subjects received payment based the outcome of their chosen lottery. The behavior of most subjects (7 out of 10) was best described by EV. This is not unexpected given the high quantitative abilities of the subjects. We also found heterogeneity in classification: One subject got classified as MVS, as identified by violations of stochastic dominance in the last few choices. 2 subjects were best described by prospect theory since they exhibited a high degree of loss aversion and risk aversion. One subject was also classified as a CRRRA-type (log-utility maximizer). Figure 2(c) shows the probability of the classified model with number of tests. Although we need a larger sample to make significant claims of the validity of different economic theories, our preliminary results indicate that subject types can be identified and there is heterogeneity in the population. They also serve as an example of the benefits of using real-time dynamic experimental design to collect data on human behavior.

6 Conclusions

In this paper, we considered the problem of adaptively selecting which noisy tests to perform in order to identify an unknown hypothesis sampled from a known prior distribution. We studied the Equivalence Class Determination problem as a means to reduce the case of noisy observations to the classic, noiseless case. We introduced EC², an adaptive greedy algorithm that is guaranteed to choose the same hypothesis as if it had observed the outcome of all tests, and incurs near-minimal expected cost among all policies with this guarantee. This is in contrast to popular heuristics that are greedy w.r.t. version space mass reduction, information gain or value of information, all of which we show can be very far from optimal. EC² works by greedily optimizing an objective tailored to differentiate between sets of observations that lead to different decisions. Our bounds rely on the fact that this objective function is adaptive submodular. We also develop EFFECXTIVE, a practical algorithm based on EC², that can be applied to arbitrary probabilistic models in which efficient exact inference is possible. We apply EFFECXTIVE to a Bayesian experimental design problem, and our results indicate its effectiveness in comparison to existing algorithms. We believe that our results provide an interesting direction towards providing a theoretical foundation for practical active learning and experimental design problems.

Acknowledgments. This research was partially supported by ONR grant N00014-09-1-1044, NSF grant CNS-0932392, NSF grant IIS-0953413, a gift by Microsoft Corporation, an Okawa Foundation Research Grant, and by the Caltech Center for the Mathematics of Information.

References

- [1] N. Balcan, A. Beygelzimer, and J. Langford. Agnostic active learning. In *ICML*, 2006.
- [2] Gowtham Bellala, Suresh Bhavnani, and Clayton Scott. Extensions of generalized binary search to group identification and exponential costs. In *Advances in Neural Information Processing Systems (NIPS)*, 2010.
- [3] Gowtham Bellala, Suresh K. Bhavnani, and Clayton D. Scott. Group-based query learning for rapid diagnosis in time-critical situations. *CoRR*, abs/0911.4511, 2009.
- [4] Colin F. Camerer. An experimental test of several generalized utility theories. *The Journal of Risk and Uncertainty*, 2(1):61–104, 1989.
- [5] V. T. Chakaravarthy, V. Pandit, S. Roy, P. Awasthi, and M. Mohania. Decision trees for entity identification: Approximation algorithms and hardness results. In *In Proceedings of the ACM- SIGMOD Symposium on Principles of Database Systems*, 2007.
- [6] K. Chaloner and I. Verdinelli. Bayesian experimental design: A review. *Statistical Science*, 10(3):273–304, Aug. 1995.
- [7] Sanjoy Dasgupta. Analysis of a greedy active learning strategy. In *NIPS*, 2004.
- [8] Sanjoy Dasgupta. Coarse sample complexity bounds for active learning. In Y. Weiss, B. Schölkopf, and J. Platt, editors, *Advances in Neural Information Processing Systems 18*, pages 235–242. MIT Press, Cambridge, MA, 2006.
- [9] Daniel Golovin and Andreas Krause. Adaptive submodularity: Theory and applications in active learning and stochastic optimization. *CoRR*, abs/1003.3967v3, 2010.
- [10] Daniel Golovin, Andreas Krause, and Debajyoti Ray. Near-optimal Bayesian active learning with noisy observations. *CoRR*, abs/1010.3091, 2010.
- [11] Andrew Guillory and Jeff Bilmes. Average-case active learning with costs. In *The 20th International Conference on Algorithmic Learning Theory*, University of Porto, Portugal, October 2009.
- [12] Giora Hanoch and Haim Levy. Efficient portfolio selection with quadratic and cubic utility. *The Journal of Business*, 43(2):181–189, 1970.
- [13] R. A. Howard. Information value theory. In *IEEE Transactions on Systems Science and Cybernetics (SSC-2)*, 1966.
- [14] D. Kahneman and A. Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–292, 1979.
- [15] S. Rao Kosaraju, Teresa M. Przytycka, and Ryan S. Borgstrom. On an optimal split tree problem. In *WADS '99: Proceedings of the 6th International Workshop on Algorithms and Data Structures*, pages 157–168, London, UK, 1999. Springer-Verlag.
- [16] Andreas Krause and Carlos Guestrin. Optimal value of information in graphical models. *Journal of Artificial Intelligence Research (JAIR)*, 35:557–591, 2009.
- [17] D. V. Lindley. On a measure of the information provided by an experiment. *Annals of Mathematical Statistics*, 27:986–1005, 1956.
- [18] D. MacKay. Information-based objective functions for active data selection. *Neural Computation*, 4(4):590–604, 1992.
- [19] Rob Nowak. Noisy generalized binary search. In Y. Bengio, D. Schuurmans, J. Lafferty, C. K. I. Williams, and A. Culotta, editors, *Advances in Neural Information Processing Systems 22*, pages 1366–1374. 2009.
- [20] John W. Pratt. Risk aversion in the small and in the large. *Econometrica*, 32(1):122–136, 1964.
- [21] D. Prelec. The probability weighting function. *Econometrica*, 66(3):497–527, 1998.
- [22] John von Neumann and Oskar Morgenstern. *Theory of Games and Economic Behaviour*. Princeton University Press, 1947.
- [23] Alice X. Zheng, Irina Rish, and Alina Beygelzimer. Efficient test selection in active diagnosis via entropy approximation. In *UAI '05, Proceedings of the 21st Conference in Uncertainty in Artificial Intelligence*, 2005.