
Near Optimal Bayesian Active Learning for Decision Making

Shervin Javdani
Carnegie Mellon University

Yuxin Chen
ETH Zürich

Amin Karbasi
ETH Zürich

Andreas Krause
ETH Zürich

J. Andrew Bagnell
Carnegie Mellon University

Siddhartha Srinivasa
Carnegie Mellon University

Abstract

How should we gather information to make effective decisions? We address Bayesian active learning and experimental design problems, where we sequentially select tests to reduce uncertainty about a set of hypotheses. Instead of minimizing uncertainty per se, we consider a set of overlapping *decision regions* of these hypotheses. Our goal is to drive uncertainty into a single decision region as quickly as possible.

We identify necessary and sufficient conditions for correctly identifying a decision region that contains all hypotheses consistent with observations. We develop a novel *Hyperedge Cutting* (HEC) algorithm for this problem, and prove that is competitive with the intractable optimal policy. Our efficient implementation of the algorithm relies on computing subsets of the complete homogeneous symmetric polynomials. Finally, we demonstrate its effectiveness on two practical applications: approximate comparison-based learning and active localization using a robot manipulator.

1 Introduction

Bayesian active learning addresses the problem of selecting a sequence of experiments, or *tests*, to determine a hypothesis consistent with observations. This fundamental problem arises in a wide range of applications such as medical procedures, content search, and robotics. It has been studied in several domains, including machine learning (Dasgupta, 2004; Balcan et al., 2006; Nowak, 2009), statistics (Lind-

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

ley, 1956; Chaloner and Verdinelli, 1995), decision theory (Howard, 1966), and others.

For instance, in *automated medical diagnosis* (Kononenko, 2001) we are presented with hypotheses about the state of a patient, and select medical tests to infer their illness. In *comparison-based learning* (Goyal et al., 2008; Karbasi et al., 2012), we infer a target in a database by sequentially presenting a user with pairs of candidates, and having the user select which is closer. In robotic *active localization*, the robot attempts to identify its own or an object’s location by probing, e.g., with touch or vision (Fox et al., 1998; Kollar and Roy, 2008; Hsiao et al., 2008; Javdani et al., 2013). In general, the goal is to gather the necessary information while minimizing test cost.

In this paper, we develop a general framework for addressing these problems. Instead of indiscriminately minimizing uncertainty about hypotheses directly, we aim to reduce uncertainty in a structured way to facilitate *decision making*. We suppose the hypothesis space is covered by a set of decision regions: Each region identifies the set of hypotheses for which it would succeed. Our goal is to select tests that quickly concentrate all consistent hypotheses in a single decision region.

Special cases of this general problem have been studied. In the so called *Optimal Decision Tree* (ODT) problem, each decision region corresponds to a single hypothesis. In this case, a greedy algorithm called *Generalized Binary Search* (GBS) is known to perform near optimally, i.e., the expected number of observations is $O(\log n)$ more than the optimum policy where n indicates the number of hypotheses (Dasgupta, 2004; Guillory and Bilmes, 2009; Kosaraju and Borgstrom, 1999). GBS greedily selects tests in expectation over the test outcomes to maximize the probability mass of eliminated hypotheses. Another special instance of our setting is the *Equivalence Class Determination* (ECD) problem (Golovin et al., 2010) where the set of hypotheses is (disjointly) partitioned— that is, decision

regions do not overlap and collectively cover the set of hypotheses. In this case, it is known that GBS performs poorly while greedily optimizing a more informative objective known as EC2 exhibits an $O(\log n)$ approximation guarantee (Golovin et al., 2010).

In both aforementioned settings, decision regions are *disjoint*. In this paper, we tackle the general case of *overlapping decision regions*, a problem that is less understood. We develop a novel surrogate objective function, which we call *Hyperedge Cutting* (HEC), and prove that the policy which greedily maximizes this objective has strong theoretical guarantees. It relies on the fact that our proposed objective function satisfies *adaptive submodularity* (Golovin and Krause, 2011), a natural diminishing returns property that generalizes the classical notion of submodularity to policies.

We empirically evaluate our algorithm on two applications: approximate comparison-based learning (Kar-basi et al., 2012), and active localization with a robot hand. In approximate comparison-based learning, a user is searching through set of items (e.g., movies), and is not particularly interested in a single item, but rather any suggestion from a given category (e.g., the horror genre). The search terminates once all items consistent with user responses are contained in a single category. Similarly, many actions in robotic manipulation, such as pushing a button or grasping an object, inherently tolerate some uncertainty. The robot need not know the exact location of an object, but rather must localize an object to a decision region to ensure it can successfully accomplish the task. An optimal policy achieves each of these with the smallest test cost.

We make the following contributions:

1. We provide a necessary and sufficient condition for identifying if a decision region contains all hypotheses that are consistent with the tests performed.
2. We develop a novel algorithm – *Hyperedge Cutting* (HEC) – and prove that it is competitive with the intractable optimal algorithm.
3. We provide an efficient way to implement our algorithm based on computing sums of the complete homogeneous symmetric polynomials.
4. We demonstrate the empirical effectiveness of our approach for both comparison-based learning and active localization in a robotic manipulation task.

2 Problem Statement

We formalize our Bayesian active learning problem by assuming a prior probability distribution P on a set of hypotheses \mathcal{H} (e.g., state of patient, location of target). By conducting tests from a set of tests \mathcal{T} , we gain infor-

mation about the true, initially unknown hypothesis.

More formally, for a given hypothesis $h \in \mathcal{H}$, running a test $t \in \mathcal{T}$ produces an outcome (deterministically) from a finite set of outcomes/observations \mathcal{O} . Thus, each hypothesis $h \in \mathcal{H}$ can be considered a function $h : \mathcal{T} \rightarrow \mathcal{O}$ mapping tests to outcomes. Suppose we have executed a set of tests $T = \{t_1, \dots, t_m\} \subseteq \mathcal{T}$ (e.g., medical tests we ran, items shown to the user, moves made by the robot), and have observed their outcomes $h(t_1), \dots, h(t_m)$. Our evidence so far is captured by a set of test-outcome pairs, $\mathcal{S} \subseteq \mathcal{T} \times \mathcal{O}$, where $\mathcal{S} = \{(t_1, h(t_1)), \dots, (t_m, h(t_m))\}$.

Upon observing \mathcal{S} , we can rule out hypotheses inconsistent with our observations. We denote the resulting set of hypotheses by

$$\mathcal{V}(\mathcal{S}) = \{h \in \mathcal{H} : \forall (t, o) \in \mathcal{S}, h(t) = o\} \quad (1)$$

In principle, we can now choose tests that reduce our uncertainty about the set of hypotheses directly. In many practical problems, we are primarily concerned about *reducing uncertainty for the purpose of making a decision*: it is not necessary to remove all uncertainty, but it is necessary to reduce uncertainty in a structured way to ensure a decision action will be successful. Choosing tests that reduce uncertainty dramatically, but still leave it unclear what action to choose, will not be effective. We now formalize this idea.

Active learning for decision making. Suppose we have a set of decisions \mathcal{R} , and the eventual goal of selecting a decision $r \in \mathcal{R}$ after gathering information. For example, in medical diagnosis, we choose a treatment; in robotic manipulation, we press a button (Fig. 5); in content search, we recommend a particular movie.

Each decision region r corresponds to the set of hypotheses for which it would succeed, i.e., $r \subseteq \mathcal{H}$. Our problem is then captured by a *hypergraph*, a generalization of a graph in which an edge can connect to any number of nodes. Briefly, a hypergraph \mathbf{G} is a pair $\mathbf{G} = (X, E)$, where X is a set of elements called *nodes*, and E is a collection of sets of X called *hyperedges*. We can specify our problem with a hypergraph, which we refer to as the *region hypergraph* $\mathbf{G}^r = (\mathcal{H}, \mathcal{R})$.¹

Note that in general, multiple decisions are equally suitable for a hypothesis: In the robot example, multiple manipulation actions may succeed for an object location (Fig. 5); in movie recommendation, the user may be indifferent among sets of movies. Hence, we allow the decision regions to overlap (Fig. 1(a)). Formally, we also assume that the set of hypotheses is covered by the collection of decision regions, i.e., $\mathcal{H} = \cup_{r \in \mathcal{R}} r$.

¹We illustrate decision regions as circles (e.g., Fig. 1(a)) - however, our method treats regions as arbitrary sets.

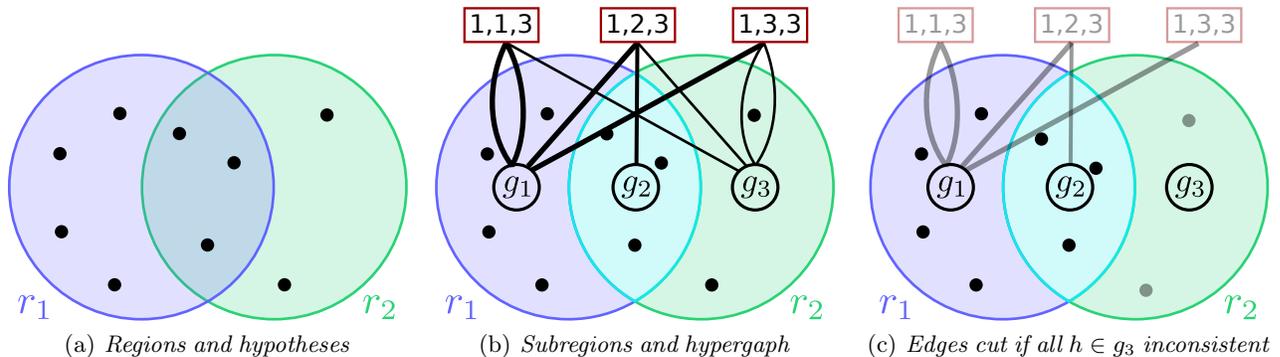


Figure 1: (a) An instance of the Decision Region Determination (DRD) problem with two decision regions. Black dots represent hypotheses and circles represent decision regions. (b) The resulting subregions and splitting hyperedges constructed by Hyperedge Cutting (HEC) algorithm. Thickness of edge represents weight, which is proportional to weight in subregion. (c) Resulting hypergraph when all hypotheses in subregion g_3 inconsistent, causing all edges to be “cut”.

The ultimate goal is to find a policy π for running tests that allows us to determine a decision region r the true hypothesis is guaranteed to lie in. In other words, upon termination we require that $\mathcal{V}(\mathcal{S}) \subseteq r$ for some $r \in \mathcal{R}$.

Thus, we seek a policy for selecting a minimal number of tests to determine a suitable decision. A *policy* π is a function from a set of evidence so far \mathcal{S} , to the next test to choose (or to stop running tests). A policy is feasible if and only if it drives all remaining uncertainty into any single decision region, $\mathcal{V}(\mathcal{S}) \subseteq r$. We define the expected cost (i.e., number of tests²) of policy π as:

$$\mathcal{C}(\pi) = \sum_{h \in \mathcal{H}} P(h) |\mathcal{T}(\pi, h)|,$$

where $\mathcal{T}(\pi, h)$ is the set of tests policy π chooses in case the correct hypothesis is h . Given this, we seek a feasible policy of minimal cost, i.e.,

$$\pi^* = \arg \min_{\pi} \mathcal{C}(\pi) \text{ s.t. } \forall h, \exists r : \mathcal{V}(\mathcal{T}(\pi, h)) \subseteq r \quad (2)$$

We call Problem (2) the *Decision Region Determination (DRD) Problem*.

Special cases of Problem (2) have been studied before. In particular, the special case where each hypothesis is contained in a dedicated region is called the Optimal Decision Tree (ODT) problem (Kosaraju and Borgstrom, 1999). More generally, the special case where the regions *partition* the hypothesis space (i.e., do not overlap), is called the Equivalence Class Determination (ECD) Problem (Golovin et al., 2010). For both of these special cases, it is known that finding a policy π for which $\mathcal{C}(\pi) \leq \mathcal{C}(\pi^*) o(\log n)$ is NP-hard (Chakaravarthy et al., 2007). Here, π^* indicates the optimum policy. To the best of our knowledge, there are no efficient algorithms with theoretical approximation guarantees for the general DRD problem. In the following, we present such an algorithm.

²Note that while we focus on tests with unit cost, our results generalize to tests with non-uniform costs.

3 The HEC Algorithm

We now introduce and analyze our algorithm – the *Hyperedge Cutting* (HEC) approach.

3.1 Overview

Our key strategy is to transform the DRD Problem (2) into an alternative representation – a different hypergraph for splitting decision regions. Observing certain test outcomes corresponds to downweighting or cutting hyperedges in this hypergraph. The construction is chosen so that cutting all hyperedges is a necessary and sufficient condition for driving all uncertainty into a single decision region. We then prove that a simple greedy algorithm, which chooses tests that reduce hyperedge weight maximally (in expectation), implements a policy that is competitive with the optimal (intractable) policy for Problem (2). In Sec. 4, we show how this greedy algorithm can be efficiently implemented.

3.2 Splitting hypergraph construction

We construct a different hypergraph, the *splitting hypergraph* \mathbf{G}^s , and define our objective on that. Here, our hyperedges are not sets, but *multisets*, a generalization of sets where members are allowed to appear more than once. As a result, a node can potentially appear in a hyperedge multiple times. The cardinality of a hyperedge refers to how many nodes it is connected to.

We observe that for solving the DRD problem, we can group together all hypotheses that share the same region assignments. We refer to this grouping as a *subregion* g , and the set of all subregions as \mathcal{G} . More formally, for any pair $h_k \in g_i$ and $h_l \in g_i$, we have $h_k \in r_j$ if and only if $h_l \in r_j$. In a slight abuse of notation, we say that a subregion is contained in a region, $g \in r$, if $\forall h \in g, h \in r$ (Fig. 1(b)). Similarly, we say that $h \in e$ if $\exists g \in e$ s.t. $h \in g$. It is easy to see that

all remaining hypotheses $\mathcal{V}(\mathcal{S})$ are contained in r if and only if all remaining subregions are contained in r .

We construct the splitting hypergraph \mathbf{G}^s over these subregions. Each subregion $g \in \mathcal{G}$ corresponds to a node. The hyperedges $e \in \mathcal{E}$ consist of all multisets of precisely k subregions, $e = \{g_1, \dots, g_k\}$, such that a single decision region does not contain them all (we will describe how k is selected momentarily). Note that hyperedges can contain the same subregion multiple times. Formally,

$$\mathcal{E} = \{e : |e| = k \wedge \nexists r \text{ s.t. } \forall h \in e, h \in r\}. \quad (3)$$

Our splitting hypergraph is defined as $\mathbf{G}^s = (\mathcal{G}, \mathcal{E})$. Fig. 1(b) illustrates the splitting hypergraph obtained from the DRD instance of Fig. 1(a).

Hyperedge Cardinality k . Key to attaining our results is the proper selection of hyperedge cardinality k . If k is too small, our results won't hold, and our algorithm won't solve the DRD problem. If k is too large, we waste computational effort, and our theoretical bounds loosen. Here, we define the cardinality we use practically. Our theorems hold for a smaller, more difficult to compute k as well. See Sec. 7 for details.

$$k = \min \left(\max_{h \in \mathcal{H}} |\{r : h \in r\}|, \max_{r \in \mathcal{R}} |\{g : g \in r\}| \right) + 1 \quad (4)$$

Note that each term is a property of the original region hypergraph \mathbf{G}^r defined in Sec. 2: $\max_h |\{r : h \in r\}|$ is the maximum degree of any node, and $\max_r |\{g : g \in r\}|$ bounds the maximum cardinality of hyperedges in \mathbf{G}^r .³

3.3 Relating DRD and HEC

How does the hypergraph capture our progress towards solving Problem (2)? Observing a set of test-outcomes $\mathcal{S} \subseteq \mathcal{T} \times \mathcal{O}$ eliminates inconsistent hypotheses, and consequently downweights or eliminates (“cuts”) incident hyperedges (Fig. 1(c)). Analogous to the definition of $\mathcal{V}(\mathcal{S})$ in (5), we define the set of hyperedges consistent with \mathcal{S} by

$$\mathcal{E}(\mathcal{S}) = \{e \in \mathcal{E} : \forall (i, o) \in \mathcal{S} \forall h \in e, h(i) = o\} \quad (5)$$

The following result guarantees that cutting all hyperedges is a necessary and sufficient condition for success, i.e., driving all uncertainty into a single decision region.

Theorem 1. *Suppose we construct a splitting hypergraph by drawing hyperedges of cardinality k according to (3). Let $\mathcal{S} \subseteq \mathcal{T} \times \mathcal{O}$ be a set of evidence. All consistent hypotheses lie in some decision region if and only if all hyperedges are cut, i.e.,*

$$\mathcal{E}(\mathcal{S}) = \emptyset \iff \exists r : \mathcal{V}(\mathcal{S}) \subseteq r$$

³It is precisely the maximum cardinality of any hyperedge if we grouped hypotheses into subregions in \mathbf{G}^r .

Thus, the DRD Problem (2) is equivalent to finding a policy of minimum cost that cuts all hyperedges. This insight suggests a natural algorithm: select tests that cut as many edges as possible (in expectation). In the following, we formalize this approach.

3.4 The Hyperedge Cutting (HEC) Algorithm

Given the above construction, we define a suitable objective function whose maximization will ensure that we pick tests to remove hyperedges quickly, thus providing us with an algorithm that identifies a correct decision region. First, we define the weight of a subregion as the sum of hypothesis weights, $p(g) = \sum_{h \in g} p(h)$. We define the weight of a hyperedge $e = \{g_1, \dots, g_k\}$ as $w(e) = \prod_{i=1}^k P(g_i)$. More generally, we define the weight of a collection of hyperedges as $w(\{e_1, \dots, e_n\}) = \sum_{l=1}^n w(e_l)$. Now, given a pair of test/observation (t, o) , we can identify the set of inconsistent hypotheses, which in turn implies the set of hyperedges that should be downweighted or removed. Formally, given a set of test/observation pairs $\mathcal{S} \subseteq \mathcal{T} \times \mathcal{O}$, we define its utility $f_{\text{HEC}}(\mathcal{S})$ as

$$f_{\text{HEC}}(\mathcal{S}) = w(\mathcal{E}) - w(\mathcal{E}(\mathcal{S})). \quad (6)$$

Thus $f_{\text{HEC}}(\mathcal{S})$ is the total mass of all the edges cut via observing set \mathcal{S} .

A natural approach to the DRD Problem is thus to seek policies that maximize (6) as quickly as possible. Arguably the simplest approach is a greedy approach that iteratively chooses the test that increases (6) as much as possible, in expectation over test outcomes.

Formally, we define the *expected marginal gain* of a test t given evidence $\mathcal{S} \subseteq \mathcal{T} \times \mathcal{O}$ as follows:

$$\Delta(t | \mathcal{S}) = \sum_h P(h | \mathcal{S}) \left(f_{\text{HEC}}(\mathcal{S} \cup \{(t, h(t))\}) - f_{\text{HEC}}(\mathcal{S}) \right)$$

Thus, $\Delta(t | \mathcal{S})$ quantifies, for test t , the expected reduction in hyperedge mass upon observing the outcome of the test. Hereby, the expectation is taken w.r.t. the distribution over hypotheses conditioned on our evidence so far. It is apparent that all hyperedges are cut if and only if $\Delta(t | \mathcal{S}) = 0$ for all tests $t \in \mathcal{T}$. Given this, our HEC Algorithm simply starts with $\mathcal{S} = \emptyset$. It then proceeds in an iterative manner, greedily selecting the test t^* that maximizes the expected marginal benefit, $t^* = \arg \max_t \Delta(t | \mathcal{S})$, observes the outcome $h(t^*)$ and adds the pair $(t^*, h(t^*))$ to \mathcal{S} . It stops as soon as all edges are cut (i.e., the marginal gain of all tests is 0).

3.5 Theoretical Analysis

The key insight behind our analysis is that the marginal gain $\Delta(t | \mathcal{S})$ satisfies two properties: *adaptive monotonicity* and *adaptive submodularity*, introduced by

Golovin and Krause (2011) and associated with certain sequential decision problems. Formally, adaptive monotonicity simply states that the benefit of each test is nonnegative, $\Delta(t \mid \mathcal{S}) \geq 0$ for all tests $t \in \mathcal{T}$ and evidence $\mathcal{S} \subseteq \mathcal{T} \times \mathcal{O}$. This is straightforward, since carrying out a test can never introduce hyperedges, but only remove them. The second, slightly more subtle property – *adaptive submodularity* – states that the marginal gain of any fixed test $t \in \mathcal{T}$ can never increase as we gain additional evidence. Formally, whenever $\mathcal{S} \subseteq \mathcal{S}' \subseteq \mathcal{T} \times \mathcal{O}$, it must hold that $\Delta(t \mid \mathcal{S}) \geq \Delta(t \mid \mathcal{S}')$. Those properties are formally established for our f_{HEC} objective and the associated marginal gain Δ in the following Theorem:

Theorem 2. *The objective function f_{HEC} defined in (6) is adaptive submodular and strongly adaptive monotone.*

Why are these properties useful? Golovin and Krause (2011) prove that for sequential decision problems satisfying adaptive monotonicity and adaptive submodularity, greedy policies are competitive with the optimal policy. In particular, as a consequence of Theorem 2 and Theorem 5.8 of Golovin and Krause (2011), we obtain the following result for our HEC Algorithm:

Theorem 3. *Assume that the prior probability distribution P on the set of hypotheses is rational. Then, the performance of π_{HEC} is bounded as follows:*

$$\mathcal{C}(\pi_{\text{HEC}}) \leq (k \ln(1/p_{\min}) + 1)\mathcal{C}(\pi^*),$$

where $p_{\min} = \min_{h \in \mathcal{H}} P(h)$.

For the special case of disjoint regions (i.e., the ECD Problem, corresponding to $k = 2$), our objective f_{HEC} is equivalent to the objective function proposed by Golovin et al. (2010), and hence our Theorem 3 strictly generalizes their result. Furthermore, in the special case where each test can have at most two outcomes, and we set $k = 1$, the HEC Algorithm is equivalent to the Generalized Binary Search algorithm for the ODT problem, and recovers its approximation guarantee.

4 Efficient Implementation

Our HEC algorithm computes $\Delta(t \mid \mathcal{S})$ for every test in \mathcal{T} , and greedily selects one at each time step. Naively computing this quantity involves constructing the splitting hypergraph \mathbf{G}^s for every possible observation, and summing the edge weights. This is computationally expensive, as constructing the graph requires enumerating every multiset of order k and checking if any region contains them all, resulting in a runtime of $O(|\mathcal{G}|^k)$. We can, however, quickly prune checks and iteratively consider multisets of growing cardinality during our computation by utilizing the following fact:

Algorithm 1 Hyperedge Weight

```

procedure HYPEREDGE WEIGHT( $\mathcal{H}, k$ )
    Compute subregions  $\mathcal{G}$  from  $\mathcal{H}$ 
     $W \leftarrow \text{CHP}_k(\mathcal{G})$ 
    Initialize queue  $Q_1$  with every subregion  $g \in \mathcal{G}$ 
    for all  $\hat{k} \leq k$  do
        for all  $\zeta_{\hat{k}} \in Q_{\hat{k}}$  do
            if  $\exists r$  s.t.  $\forall h \in \zeta_{\hat{k}}, h \in r$  then
                 $W \leftarrow W - \prod_{g \in \zeta_{\hat{k}}} p(g) \text{CHP}_{k-\hat{k}}(\zeta_{\hat{k}})$ 
                Add all supersets of  $\zeta_{\hat{k}}$  to  $Q_{\hat{k}+1}$ 
    return  $W$ 
    
```

Proposition 1. *A set of subregions G shares a region only if all subsets $G' \subset G$ also share that region.*

4.1 Utilizing Complete Homogeneous Symmetric Polynomials

Our general strategy will be to compute the sum of weights over *all* multisets of cardinality k , and subtract those that correspond to a shared region. To do so efficiently, we identify algebraic structure in computing a sum of multisets, where a multiset corresponds to a product. Namely, it is equivalent to computing a complete homogeneous symmetric polynomial.

For any $G \subseteq \mathcal{G}$ and cardinality \hat{k} , we define $\mathcal{G}_{\hat{k}}(G)$ as all multisets over groups G of cardinality \hat{k} . Unlike hyperedges, these multisets can share a region. Formally

$$\mathcal{G}_{\hat{k}}(G) = \{\{g_1, \dots, g_{\hat{k}}\} \subseteq G\}$$

Recall that $w(\mathcal{G}_{\hat{k}}(G)) = \sum_{g \in \mathcal{G}_{\hat{k}}(G)} \prod_g P(g)$. Computing $w(\mathcal{G}_{\hat{k}}(G))$ can be performed efficiently as this quantity is exactly equivalent to the *complete homogeneous symmetric polynomial* (CHP) of degree \hat{k} over G . We will briefly review a well known variant of the Newton-Girard formulae which will make an efficient algorithm for computing $w(\mathcal{G}_{\hat{k}}(G))$ clear.

Define any set of variables $\mathbf{x} = \{x_1, \dots, x_n\}$.

$$PS_i(\mathbf{x}) = \sum_{x \in \mathbf{x}} x^i$$

$$CHP_i(\mathbf{x}) = \sum_{l_1 + \dots + l_n = i; l_j \geq 0} \prod_{x_j \in \mathbf{x}} x_j^{l_j}$$

Here PS_i is the i -th *power sum*, and CHP_i is the i -th complete homogeneous symmetric polynomial.

We have the identity (Macdonald, 1998; Seroul, 2000):

$$CHP_i(\mathbf{x}) = \frac{1}{i} \sum_{j=1}^i CHP_{i-j}(\mathbf{x}) PS_j(\mathbf{x})$$

Thus, we iteratively compute $CHP_1(G) \dots CHP_{\hat{k}}(G)$ to compute $w(\mathcal{G}_{\hat{k}}(G)) = CHP_{\hat{k}}(G)$ with runtime $O(\hat{k}|G|)$.

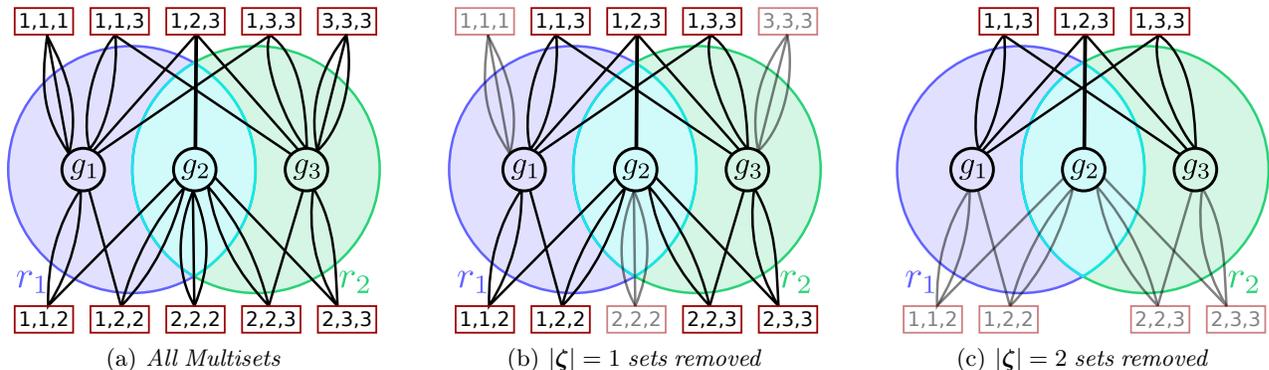


Figure 2: A depiction of our algorithm as hyperedges. (a) The equivalent hyperedges of $CHP_3(\mathcal{G})$. (b) First iteration of Alg. 1 which removes all $|\zeta| = 1$ (light edges) by subtracting $g_1CHP_2(\{g_1\}) + g_2CHP_2(\{g_2\}) + g_3CHP_2(\{g_3\})$. (c) Second iteration of Alg. 1 which removes all $|\zeta| = 2$ (light edges) by subtracting $g_1g_2CHP_1(\{g_1, g_2\}) + g_2g_3CHP_1(\{g_2, g_3\})$

We now turn our attention to efficiently computing the weight of all multisets that correspond to subregions encapsulated by a region. Let ζ be a set (not multiset) of subregions that shares a region. Formally:

$$\zeta = \{g_1 \dots g_{\hat{k}}\} \quad \hat{k} \leq k, \nexists r \text{ s.t. } \zeta \subseteq r$$

We compute the term corresponding to ζ we subtract from $CHP_k(\mathcal{G})$ (weight of all multisets), as ζ shares a region. To avoid double counting, we force the term to include $\prod_{g \in \zeta} p(g)$ as a factor, i.e. if we think of a hyperedge as a product, we force one link to each element of ζ .

$$\begin{aligned} w(\zeta) &= \prod_{g \in \zeta} p(g) \sum_{l_1 + \dots + l_{\hat{k}} = k - \hat{k}; l_i > 0} p(g_1)^{l_1} \dots p(g_{\hat{k}})^{l_{\hat{k}}} \\ &= \prod_{g \in \zeta} p(g) CHP_{k - \hat{k}}(\zeta) \end{aligned}$$

Using this, we compute $w(\mathcal{E}) = CHP_k(\mathcal{G}) - \sum_{\zeta \subseteq \mathcal{G}} w(\zeta)$ by finding every set $\zeta \subseteq \mathcal{G}$ that shares a region. Furthermore, we can utilize Proposition 1 to prune sets, and only consider $\zeta_{\hat{k}+1}$ which are supersets of any $\zeta_{\hat{k}}$. The algorithm is detailed in Alg. 1, and depicted in Fig. 2.

Additionally, we note that region assignments do not change as observations are received. In practice, we find all sets of subregions that share a region once. At each time step, we need only sum over the terms corresponding to remaining hypotheses.

Note that in the worst case, this algorithm still has complexity $O(|\mathcal{G}|^k)$. This occurs when many, at least k , subregions share a single region. The complexity is then controlled by how many distinct subregions a single region can be shattered into, and the largest number of regions a single hypothesis can belong to. However, for many practical problems, we might expect many regions to be separated, e.g., when $|\mathcal{R}| \gg k$. In this case, Alg. 1 will be significantly more efficient.

Finally, we note that we can utilize an *accelerated adaptive greedy* algorithm applicable to all adaptive

submodular functions, which directly uses the diminishing returns property to skip reevaluation of actions (Golovin and Krause, 2011).

5 Experiments

In this section, we empirically evaluate HEC on the two applications - approximate comparison-based learning and touch based localization with a robotic end effector.

We compare HEC with five baselines. The first two are variants of algorithms for the specialized versions of the DRD problem described earlier - generalized binary search (Nowak, 2009) and equivalence class edge cutting (Golovin et al., 2010). For generalized binary search (GBS), we assign each hypothesis to its own decision region, and run HEC on this hypothesis-region assignment until only one hypothesis remains. To apply equivalence class edge cutting (EC2), decision regions must be disjoint. Thus, we randomly assign each hypothesis to one of the decision regions that it belongs to, and run EC2 until only one of these new regions remains. For each of these, we also run a slightly modified version, termed GBS-HEC and EC2-HEC respectively, which selects tests based on these algorithms, but terminates once all hypotheses are contained in one decision region in the original DRD problem (i.e. when the HEC termination condition is met).

The last baseline is a classic heuristic from decision theory: myopic value of information (VoI) (Howard, 1966). We define a utility function $U(h, r)$ which is 1 if $h \in r$ and 0 otherwise. The utility of $\mathcal{V}(\mathcal{S})$ corresponds to the maximum expected utility of any decision region, i.e., the expected utility if we made a decision now. VoI greedily chooses the test that maximizes (in expectation over observations) the gain in this utility. Note that if we could solve the intractable problem of nonmyopically optimizing VoI (i.e., look ahead arbitrarily to consider outcomes of sequences of tests),

we could solve the DRD problem optimally. In some sense, HEC can be viewed as a surrogate function for nonmyopic value of information.

5.1 Approximate comparison-based learning

We evaluate HEC on the *MovieLens 100k*⁴ dataset, which consists of 1 to 5 ratings of 1682 movies from 943 users. We partition movies into decision regions using these ratings, with the goal of recommending any movie in a decision region. In order to get a similarity measurement between movies, we map them into a 10-dimensional feature space by computing a low-rank approximation of the user/rating matrix through SVD. We then use k -means to partition the set of movies into $|\mathcal{R}|$ (non-overlapping) clusters, corresponding to decision regions. Each movie is then assigned to the α closest cluster centroids. See Fig. 4 for an illustration. A test corresponds to comparing two movies, an observation to selecting one of the two, and consist hypotheses are those which are closer to the selected movie (euclidean distance in 10-dimensional feature space).

Each experiment corresponds to sampling one movie as the “true” movie. The size of a decision region determines how close our solution is to this (exact) target hypothesis. As the number of regions increases, the size of each decision region shrinks. As a result, the problem requires the selected movie be closer to the true target, at the expense of increased query complexity. Fig. 3(a) shows the query complexity of different algorithms as a function of the number of regions, with the cardinality of the HEC hypergraph fixed to $k = 3$ (i.e., each hypothesis belongs to two decision regions). An extreme case is when there are only two regions and all hypotheses belong to both regions, giving a query complexity of 0. Other than that, we see that HEC performs consistently better than other methods (e.g., to identify the true region out of 8 regions, it takes on average 6.7 queries for HEC, as opposed to 8 queries for EC2-HEC, 8.5 queries for GBS-HEC, and 10.3 queries for VoI).

To see how the cardinality and region overlap influence performance, we compare the query complexity of different algorithms by varying the number of regions each hypothesis is assigned to. If we assign more regions to a hypothesis, then the search result is allowed to be further away from the true target, and thus the number of queries required for approximated search should be smaller. Fig. 3(b) demonstrates such an effect. We fix the number of clusters to 12, and vary the number of assigned regions (and thus the hyperedge cardinality) from 1 to 4 (k from 2 to 5, respectively). We see that higher cardinality enables HEC to save more queries. For $k = 5$, it takes HEC 5.3 queries to identify a movie,

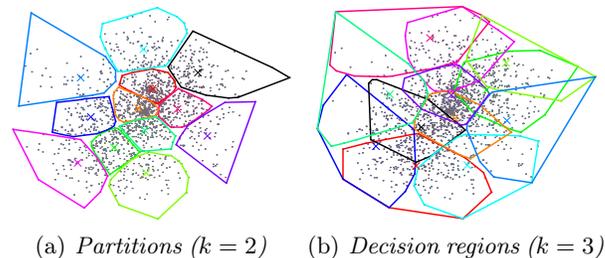


Figure 4: A 2-d illustration of (overlapping) decision regions for *MovieLens 100k* experiments. Dots represent movies, cross markers represent cluster centroids, and colored polygons represent decision region boundaries. (a) Movies are partitioned into 12 disjoint clusters. (b) Each movie is assigned to the two closest centroids.

k	2	3	4	5
t(HEC)	0.026s	0.071s	2.5s	< 2min

Table 1: Running time of HEC on *MovieLens 100k* with different cardinality k ($|\mathcal{R}| = 12$)

whereas VoI, GBS-HEC, and EC2-HEC took 8.8, 7.4, and 6.4 queries, respectively. Additionally, Table 5.1 shows the running time of HEC for these instances. We see that the accelerated implementation described in Sec. 4 enables HEC to run efficiently with reasonable hyperedge cardinality on this data set.

5.2 Touch Based localization

We evaluate HEC on a simple robotic manipulation example. Our task is to push a button with the finger of a robotic end effector. Given a distribution over object location, we generate a set of decisions, corresponding to the end effector going to a particular pose and moving forward in a straight line. Each of these decisions will succeed on a subset of hypotheses, corresponding to a decision region. Decision regions may overlap, as a button can be pushed with many decision actions. See Fig. 5.

All hypotheses are not contained in a single decision region, so we perform tests to reduce uncertainty. These tests correspond to *guarded moves* (Will and Grossman, 1975), where the end effector moves along a path until contact is sensed. After sensing contact, hypotheses are updated by eliminating object locations which could not have produced contact, e.g., if they are far away. Our goal is to find the shortest sequence of tests such that after performing them, there is a single button-push decision that would succeed for all remaining hypotheses.

Given some object location X_s , we generate an initial set of 2000 hypotheses \mathcal{H} by sampling from $N(\mu, \Sigma)$ with $\mu = X_s$, and Σ a diagonal matrix with $\Sigma_{xx} = \Sigma_{yy} = \Sigma_{zz} = 0.04$. The robot generates 50

⁴<http://www.grouplens.org/datasets/movielens/>

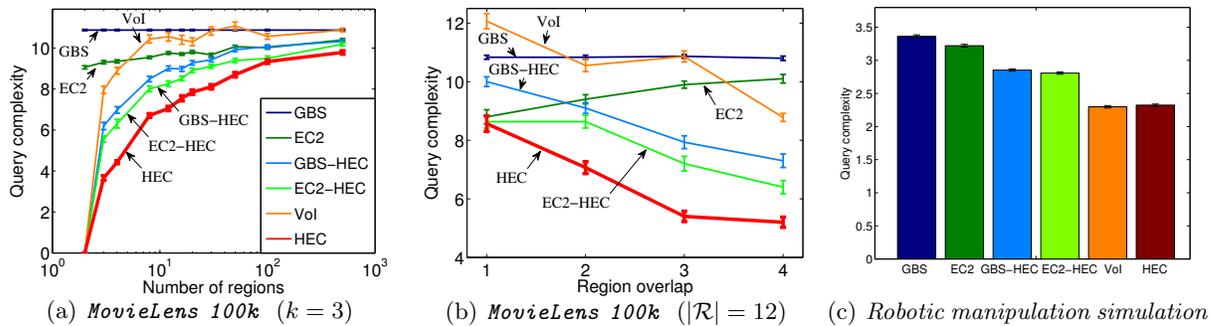


Figure 3: Results on *MovieLens 100k* and *Robot* experiments. (a) Performance as we vary the number of regions $|\mathcal{R}|$. (b) Performance as we vary the cardinality k . (c) Average performance of different algorithms across button push instances.

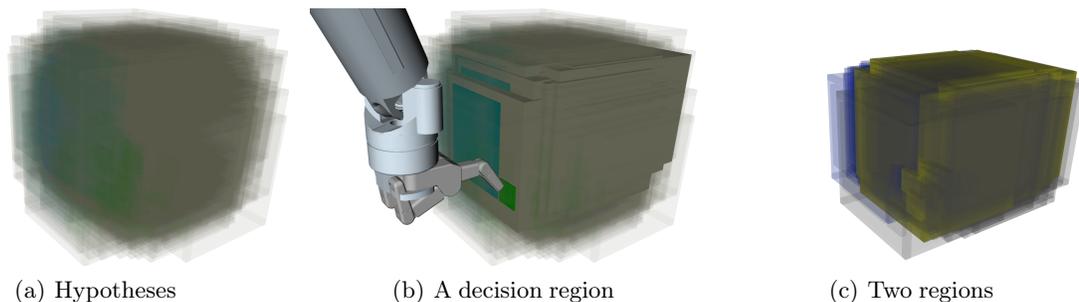


Figure 5: Touch based localization for pushing the button of a microwave. Given hypotheses over object location (a), decision actions are generated. The corresponding decision regions are computed by forward simulating to find hypotheses for which it would succeed (b). Decision regions will overlap. In (c), we see two regions (blue and grey) and their overlap (yellow).

decision regions by picking different locations and simulating the end effector forward, and noting which object poses it would succeed on. Hypotheses range from being in zero decision regions to 6, giving us a cardinality $k = 7$. For tests, the robot generates 150 guarded moves by sampling a random start location and orientation.

We conduct experiments on 10 random environments, and randomly sample 100 hypotheses to be the “true” object location (for producing observations during execution), for a total of 1000 experiments. Fig. 3(c) shows the query complexity of different algorithms averaged over these instances. We see that HEC performs well, outperforming GBS, GBS-HEC, EC2, and EC2-HEC handily. Note that myopic VoI performs essentially the same as HEC on these experiments. This is likely due to the short horizon, where 2-3 actions were usually sufficient for reducing uncertainty to a single decision region. We would expect that for longer horizons, myopic VoI would not perform as well.

6 Conclusions

In this paper, we have addressed the problem of active learning in order to facilitate decision making. We defined the Decision Region Determination (DRD) problem, requiring that at the end of information gathering, all remaining hypotheses are confined within a single

decision region (i.e., do not require further distinction from a decision making point of view). To address this problem, we proposed an equivalent representation in terms of a hypergraph. We prove that eliminating all edges in this hypergraph is a necessary and sufficient condition for success, suggesting a natural objective function. We show that this objective satisfies adaptive monotonicity and adaptive submodularity. This insight enabled us to prove that a greedy policy for removing hyperedges (HEC) has an approximation guarantee compared to the optimal policy. Finally, we note that at each iteration, we compute a particular polynomial, and can utilize a faster algorithm through efficient computations of complete homogeneous symmetric polynomials.

While our algorithm enables us to tackle problems of reasonable size, our computation is still exponential in hyperedge cardinality k . Additionally, our current scheme assumes *noise-free* observations, where a hypothesis deterministically maps a test to an observation. We hope to alleviate these limitations in future work.

Acknowledgements

This work was supported in part by the Intel Embedded Computing ISTC, NSF Grant No. 0946825, NSF-IIS-1227495, DARPA MSEE FA8650-11-1-7156, ERC StG 307036, and a Microsoft Research Faculty Fellowship.

References

- N. Balcan, A. Beygelzimer, and J. Langford. Agnostic active learning. In *ICML*, 2006.
- V. T. Chakaravarthy, V. Pandit, S. Roy, P. Awasthi, and M. K. Mohania. Decision trees for entity identification: Approximation algorithms and hardness results. In *PODS*, 2007.
- K. Chaloner and I. Verdinelli. Bayesian experimental design: A review. *Statistical Science*, 10:273–304, 1995.
- S. Dasgupta. Analysis of a greedy active learning strategy. In *NIPS*, 2004.
- D. Fox, W. Burgard, and S. Thrun. Active markov localization for mobile robots. *Robotics and Autonomous Systems*, 25:195–207, 1998.
- D. Golovin and A. Krause. Adaptive submodularity: Theory and applications in active learning and stochastic optimization. *JAIR*, 42(1):427–486, 2011.
- D. Golovin, A. Krause, and D. Ray. Near-optimal bayesian active learning with noisy observations. In *NIPS*, 2010.
- N. Goyal, Y. Lifshits, and H. Schütze. Disorder inequality: A combinatorial approach to nearest neighbor search. In *WSDM*, 2008.
- A. Guillory and J. Bilmes. Average-case active learning with costs. In *ALT*, 2009.
- R. A. Howard. Information value theory. In *IEEE Trans. Syst. Sci. Cybernetics*, 1966.
- K. Hsiao, T. Lozano-Pérez, and L. P. Kaelbling. Robust belief-based execution of manipulation programs. In *WAFR*, 2008.
- S. Javdani, M. Klingensmith, J. A. D. Bagnell, N. Pollard, and S. Srinivasa. Efficient touch based localization through submodularity. In *IEEE ICRA*, 2013.
- A. Karbasi, S. Ioannidis, and L. Massoulié. Comparison-based learning with rank nets. In *ICML*, 2012.
- T. Kollar and N. Roy. Efficient optimization of information-theoretic exploration in slam. In *AAAI*, 2008.
- I. Kononenko. Machine learning for medical diagnosis: history, state of the art and perspective. *Artif. Intell. Med.*, 23:89–109, 2001.
- R. S. Kosaraju and T. M. P. A. R. S. Borgstrom. On an optimal split tree problem. In *WADS*, 1999.
- D. V. Lindley. On a measure of the information provided by an experiment. *Ann. Math. Stat.*, 27:986–1005, 1956.
- I. G. Macdonald. *Symmetric Functions and Hall Polynomials*. Oxford mathematical monographs. Clarendon Press, 1998. ISBN 9780198504504.
- R. Nowak. Noisy generalized binary search. In *NIPS*, 2009.
- R. Seroul. *Programming for Mathematicians*. Universitext - Springer-Verlag. Springer, 2000. ISBN 9783540664222.
- P. M. Will and D. D. Grossman. An experimental system for computer controlled mechanical assembly. *IEEE Trans. Computers*, 24(9):879–888, 1975.

7 Appendix

In this section, we provide proofs for theorems stated throughout the paper.

7.1 k for bounds

We start by showing that for a properly defined k , the DRD problem is solved ($\mathcal{V}(\mathcal{S}) \subseteq r$) if and only if the HEC objective is maximized. However, we sometimes require a slightly greater k to ensure the objective f_{HEC} is adaptive submodular. We define these below.

Let R be a set of regions, the length of which is related to k . To get equivalence of the DRD and HEC, we require that for every region in R , there is some hypothesis in all but one region of R .

$$R_{\text{iff}} = \arg \max_R |R| \quad \text{s.t. } \forall r \in R, \exists h : h \notin r, h \in R \setminus r$$

$$k_{\text{iff}} = |R_{\text{iff}}|$$

Sometimes, this is not sufficient for adaptive submodularity. For this, we also require that there is some hypothesis in every region of R , and we also add one to the length of R .

$$R_{\text{as}} = \arg \max_R |R| \quad \text{s.t. } \textcircled{1} \exists \tilde{h} \in R \quad \textcircled{2} \forall r \in R, \exists h : h \notin r, h \in R \setminus r$$

$$k_{\text{as}} = |R_{\text{as}}| + 1$$

Before moving on, we prove that $k_{\text{as}} \geq k_{\text{iff}}$.

Proposition 2. $k_{\text{as}} \geq k_{\text{iff}}$

Proof. There are two cases:

1. $\exists h \in R_{\text{iff}}$. In this case, $R_{\text{as}} = R_{\text{iff}}$ and $k_{\text{as}} = |R_{\text{as}}| + 1 = k_{\text{iff}} + 1$.
2. $\nexists h \in R_{\text{iff}}$. Define $\tilde{R} = R_{\text{iff}} \setminus r$ for some $r \in R_{\text{iff}}$. We know by definition of R_{iff} that $\exists h \in \tilde{R}$. Additionally, we know by definition of k_{iff} that $\forall r \in \tilde{R}, \exists h, h \notin r, h \in R_{\text{iff}} \setminus r$, so it follows that $h \in \tilde{R} \setminus r$. Therefore, we know \tilde{R} satisfies the constraints for R_{as} , and $k_{\text{as}} \geq |\tilde{R}| + 1 = |R_{\text{iff}}| = k_{\text{iff}}$. □

Our algorithm actually utilizes $k = \min \left(\max_{h \in \mathcal{H}} |\{r : h \in r\}|, \max_{r \in \mathcal{R}} |\{g : g \in r\}| \right) + 1$. We briefly show that each of these also upper bound k_{as} .

Proposition 3. $\max_{h \in \mathcal{H}} |\{r : h \in r\}| + 1 \geq k_{\text{as}}$

Proof. Note that condition $\textcircled{1}$ in R_{as} bounds $|R_{\text{as}}|$ by $\max_{h \in \mathcal{H}} |\{r : h \in r\}|$. The result follows. □

Proposition 4. $\max_{r \in \mathcal{R}} |\{g : g \in r\}| + 1 \geq k_{\text{as}}$

Proof. Let r be an element of R_{as} . By definition, it is required that at least $|R_{\text{as}}|$ different subregions $g_1 \cdots g_{|R_{\text{as}}|}$ be in that region - one which is in every other region in R_{as} to satisfy condition $\textcircled{1}$, and $|R_{\text{as}}| - 1$ which are in all but one of the $R_{\text{as}} - 1$ other regions to satisfy condition $\textcircled{2}$. The result follows. □

Thus, we can utilize $k = \min \left(\max_{h \in \mathcal{H}} |\{r : h \in r\}|, \max_{r \in \mathcal{R}} |\{g : g \in r\}| \right) + 1$ and apply the proofs using cardinality at least k_{as} and k_{iff} . While our bounds and algorithm are better if we knew the correct k_{as} to use, finding that value is itself hard to compute - thus, our implementation uses the value defined in Sec. 3 and copied above.

7.2 Theorem 1: Equivalence of DRD and HEC

Proof. We first prove that if all h are contained in one region, then all edges are cut, i.e. $\exists r : \mathcal{V}(\mathcal{S}) \subseteq r \Rightarrow \mathcal{E}(\mathcal{S}) = \emptyset$. This is by construction, since a hyperedge $e \in \mathcal{E}(\mathcal{S})$ is only between subregions (or hypotheses) that do not share any regions. More concretely, our definition of e requires $\exists r$ s.t. $\forall h \in e : h \in r$. Since all remaining nodes $\mathcal{V}(\mathcal{S}) \subseteq r$, there will be no such set of hypotheses.

Next, we prove that if all edges are removed, then all h are contained in one region, i.e., $\mathcal{E}(\mathcal{S}) = \emptyset \Rightarrow \exists r : \mathcal{V}(\mathcal{S}) \subseteq r$. Clearly, if we set $|\mathcal{V}(\mathcal{S})| \leq k$, this condition would be met - $\mathcal{E}(\mathcal{S})$ would check every subset of $\mathcal{V}(\mathcal{S})$ to see if they shared a region, and would draw a hyperedge i.f.f. they do not. To complete the proof, we will make use of the following lemma:

Lemma 1. Define β as some constant s.t. $\beta \geq k$. $\forall H \subseteq \mathcal{H}, |H| = \beta, \exists r : H \subseteq r \Rightarrow \forall \{H \cup h\} \subseteq \mathcal{H}, \exists r : \{H \cup h\} \in r$

Proof. For the sake of contradiction, suppose $\nexists r : \{H \cup h\} \in r$. This must mean $h \notin H$. Let $\{H \cup h\} = \{h_1, h_2, \dots, h_{\beta+1}\}$. Let H_i be the subset of $\{H \cup h\}$ which does not include the i th h from $\{H \cup h\}$, i.e. $H_i = \{h_1, \dots, h_{i-1}, h_{i+1}, \dots, h_{\beta+1}\}$. By assumption, we know $\exists r : H_i \in r$. Let r_i be that region for H_i . If $r_i = r_j$, for any i, j , this would imply $\{H_i \cup H_j\} = \{H \cup h\} \in r_i$. Thus, each r_i must be unique if $\exists r : \{H \cup h\} \in r$. Furthermore, this implies $h_i \notin r_i$, and $h \in r_j, \forall j \neq i$. Let $R_{\beta+1} = \{r_1 \dots r_{\beta+1}\}$. By definition of β , we know $\beta \geq k \geq k_{\text{iff}}$. But this causes a contradiction - by definition of k_{iff} , the maximum set of regions R where $h_i \notin r_i, h_i \in r_j, \forall j \neq i$ is k_{iff} . But $R_{\beta+1}$ would require such a set of regions where $|R_{\beta+1}| = \beta + 1 \geq k_{\text{iff}} + 1$. Thus, we have a contradiction, and have shown $\exists r : \{H \cup h\} \in r$. \square

By construction, we know that if $\mathcal{E}(\mathcal{S}) = \emptyset \Rightarrow \forall H \subseteq \mathcal{H}, |H| \leq k, \exists r : H \subseteq r$. Applying Lemma 1 inductively, this implies, $\forall \{H \cup h_1\} \subseteq \mathcal{V}(\mathcal{S}), \exists r : \{H \cup h_1\} \subseteq r \Rightarrow \forall \{H \cup h_1 \cup h_2\} \subseteq \mathcal{V}(\mathcal{S}), \exists r : \{H \cup h_1 \cup h_2\} \subseteq r \Rightarrow \dots \Rightarrow \exists r : \mathcal{V}(\mathcal{S}) \subseteq r$. \square

7.3 Theorem 2: strong adaptive monotonicity and adaptive submodularity

Proof. We start with showing our formulation is strongly adaptive monotone.

Lemma 2. The function f_{HEC} described above is strongly adaptive monotone, i.e.

$$f_{\text{HEC}}(\mathcal{S} \cup \{(t, h(t))\}) - f_{\text{HEC}}(\mathcal{S}) \geq 0 \quad \forall t, h$$

Proof. This states that our utility function must always increase as we take additional actions and receive observations. Intuitively, we can see that additional action observation pairs can only cut edges, and thus our utility function always increases. More concretely:

$$\begin{aligned} & f_{\text{HEC}}(\mathcal{S} \cup \{(t, h(t))\}) - f_{\text{HEC}}(\mathcal{S}) \\ &= (w(\mathcal{E}) - w(\mathcal{E}(\mathcal{S} \cup \{(t, h(t))\}))) - (w(\mathcal{E}) - w(\mathcal{E}(\mathcal{S}))) \\ &= w(\mathcal{E}(\mathcal{S})) - w(\mathcal{E}(\mathcal{S} \cup \{(t, h(t))\})) \\ &= w(\{e \in \mathcal{E} : \forall (i, o) \in \mathcal{S} \forall \tilde{h} \in e, \tilde{h}(i) = o\}) \\ &\quad - w(\{e \in \mathcal{E} : \forall (i, o) \in \mathcal{S} \forall \tilde{h} \in e, \tilde{h}(i) = o, \tilde{h}(t) = h(t)\}) \quad \text{by definition of } \mathcal{E}(\mathcal{S}) \\ &= w(\{e \in \mathcal{E} : \forall (i, o) \in \mathcal{S} \forall \tilde{h} \in e, \tilde{h}(i) = o, \tilde{h}(t) \neq h(t)\}) \\ &\geq 0 \quad \text{since } w(e) \geq 0 \quad \forall e \end{aligned}$$

\square

Next, we prove that our formulation is adaptive submodular:

Lemma 3. The function f_{HEC} described above is adaptive submodular for any prior with rational values, i.e. for $\mathcal{S} \subseteq \hat{\mathcal{S}} \subseteq \mathcal{T} \times \mathcal{O}$

$$\Delta(t|\mathcal{S}) \geq \Delta(t|\hat{\mathcal{S}}) \quad \forall t \in \mathcal{T} \setminus \mathcal{S}_T$$

where \mathcal{S}_T are the set of tests in \mathcal{S} .

Proof. This states that our expected utility for a fixed action t decreases as we take additional actions and receive observations. We rewrite our expected marginal utility in a more convenient form:

$$\begin{aligned}\Delta(t|\mathcal{S}) &= \sum_h P(h|\mathcal{S}) \left(f_{\text{HEC}}(\mathcal{S} \cup \{(t, h(t))\}) - f_{\text{HEC}}(\mathcal{S}) \right) \\ &= \sum_h P(h|\mathcal{S}) \left([w(\mathcal{E}) - w(\mathcal{E}(\mathcal{S} \cup \{(t, h(t))\}))] - [w(\mathcal{E}) - w(\mathcal{E}(\mathcal{S}))] \right) \\ &= \sum_h P(h|\mathcal{S}) \left(w(\mathcal{E}(\mathcal{S})) - w(\mathcal{E}(\mathcal{S} \cup \{(t, h(t))\})) \right)\end{aligned}$$

For convenience, we define n_i^o to be the total probability mass in g_i consistent with all evidence in \mathcal{S} and observation o . We define n_i and n^o similarly. More formally:

$$\begin{aligned}n_i^o &= \sum_{h \in g_i} P(h) \mathbb{1}(h \in \mathcal{V}(\mathcal{S} \cup \{(t, o)\})) \\ n_i &= \sum_{o \in \mathcal{O}} n_i^o \\ n^o &= \sum_{g_i \in \mathcal{G}} n_i^o \\ N &= \sum_{g_i \in \mathcal{G}} \sum_{o \in \mathcal{O}} n_i^o \\ w(\mathcal{E}(\mathcal{S})) &= \sum_{e \in \mathcal{E}} \prod_{i \in e} n_i\end{aligned}$$

Similarly, we can also write $w(\mathcal{E}(\mathcal{S} \cup \{(t, o)\})) = \sum_{e \in \mathcal{E}} \prod_{i \in e} n_i^o$. We can rewrite our objective as:

$$\begin{aligned}\Delta(t|\mathcal{S}) &= \sum_h P(h|\mathcal{S}) \left(\sum_{e \in \mathcal{E}} \prod_{i \in e} n_i - \sum_{e \in \mathcal{E}} \prod_{i \in e} n_i^{h(t)} \right) \\ &= \sum_o \frac{n^o}{N} \left(\sum_{e \in \mathcal{E}} \prod_{i \in e} n_i - \sum_{e \in \mathcal{E}} \prod_{i \in e} n_i^o \right) \\ &= \sum_{e \in \mathcal{E}} \prod_{i \in e} n_i - \sum_o \frac{n^o}{N} \sum_{e \in \mathcal{E}} \prod_{i \in e} n_i^o\end{aligned}$$

Similarly, we define variables for the evidence $\hat{\mathcal{S}}$, i.e. \hat{n}_i^o for the total probability mass in g_i consistent with all evidence in $\hat{\mathcal{S}}$ and observation o :

$$\Delta(t|\hat{\mathcal{S}}) = \sum_{e \in \mathcal{E}} \prod_{i \in e} \hat{n}_i - \sum_o \frac{\hat{n}^o}{\hat{N}} \sum_{e \in \mathcal{E}} \prod_{i \in e} \hat{n}_i^o$$

We rewrite what we would like to show as:

$$\begin{aligned}&\Delta(t|\mathcal{S}) - \Delta(t|\hat{\mathcal{S}}) \\ &= \left(\sum_{e \in \mathcal{E}} \prod_{i \in e} n_i - \sum_o \frac{n^o}{N} \sum_{e \in \mathcal{E}} \prod_{i \in e} n_i^o \right) - \left(\sum_{e \in \mathcal{E}} \prod_{i \in e} \hat{n}_i - \sum_o \frac{\hat{n}^o}{\hat{N}} \sum_{e \in \mathcal{E}} \prod_{i \in e} \hat{n}_i^o \right) \\ &\geq 0\end{aligned}$$

We will show that for any single action observation pair, which corresponds to eliminating a single hypothesis, the expected utility of a test will always decrease. General adaptive submodularity, which states the expected utility decreases with any additional evidence, follows easily. For convenience, we consider rescaling our function so that all n_i^o are integers, which is possible since we assumed a rational prior. Note that a function f is adaptive submodular i.f.f. cf is adaptive submodular for any constant $c > 0$, so showing adaptive submodularity in the rescaled setting implies adaptive submodularity for our setting.

Lemma 3.1. *If we remove one hypothesis from subregion k which agrees with observation c , i.e.*

$$\widehat{n}_i^o = \begin{cases} n_i^o - 1 & \text{if } i = l \text{ and } o = c \\ n_i^o & \text{else} \end{cases}$$

then

$$\Delta = \left(\sum_{e \in \mathcal{E}} \prod_{i \in e} n_i - \sum_o \frac{n^o}{N} \sum_{e \in \mathcal{E}} \prod_{i \in e} n_i^o \right) - \left(\sum_{e \in \mathcal{E}} \prod_{i \in e} \widehat{n}_i - \sum_o \frac{\widehat{n}^o}{\widehat{N}} \sum_{e \in \mathcal{E}} \prod_{i \in e} \widehat{n}_i^o \right) \geq 0$$

Proof. Based on our definitions, it follows that:

$$\begin{aligned} \widehat{n}_i &= \begin{cases} n_i - 1 & \text{if } i = l \\ n_i & \text{else} \end{cases} \\ \widehat{n}^o &= \begin{cases} n^o - 1 & \text{if } o = c \\ n^o & \text{else} \end{cases} \\ \widehat{N} &= N - 1 \end{aligned}$$

We split the difference into three terms:

$$\begin{aligned} \Delta^a &= \sum_{e \in \mathcal{E}} \left(\prod_{i \in e} n_i - \prod_{i \in e} \widehat{n}_i \right) \\ \Delta^b &= \sum_{o \in \mathcal{O} \setminus c} \sum_{e \in \mathcal{E}} \left(-\frac{n^o}{N} \prod_{i \in e} n_i^o + \frac{\widehat{n}^o}{\widehat{N}} \prod_{i \in e} \widehat{n}_i^o \right) \\ \Delta^c &= \sum_{e \in \mathcal{E}} \left(-\frac{n^c}{N} \prod_{i \in e} n_i^c + \frac{\widehat{n}^c}{\widehat{N}} \prod_{i \in e} \widehat{n}_i^c \right) \\ \Delta^a + \Delta^b + \Delta^c &= \Delta \end{aligned}$$

To aid in notation, we define $\mathcal{E}_l = \{e \in \mathcal{E} : g_l \in e\}$, hyperedges that contain region l , and $\overline{\mathcal{E}}_l = \mathcal{E} \setminus \mathcal{E}_l$, all other hyperedges. Additionally, let $|e_l|$ be the number of times g_l appears in the multiset e .

First term:

$$\begin{aligned} \Delta^a &= \sum_{e \in \mathcal{E}} \left(\prod_{i \in e} n_i - \prod_{i \in e} \widehat{n}_i \right) \\ &= \sum_{e \in \overline{\mathcal{E}}_l} \left[\prod_{i \in e} n_i - \prod_{i \in e} \widehat{n}_i \right] + \sum_{e \in \mathcal{E}_l} \left[\prod_{i \in e} n_i - \prod_{i \in e} \widehat{n}_i \right] \\ &= \sum_{e \in \overline{\mathcal{E}}_l} \left[\prod_{i \in e} n_i - \prod_{i \in e} n_i \right] + \sum_{e \in \mathcal{E}_l} \left[\left(\prod_{i \in e, i \neq l} n_i \right) n_l^{|e_l|} - \left(\prod_{i \in e, i \neq l} n_i \right) (n_l - 1)^{|e_l|} \right] \\ &= \sum_{e \in \mathcal{E}_l} \left(\prod_{i \in e, i \neq l} n_i \right) \left(n_l^{|e_l|} - (n_l - 1)^{|e_l|} \right) \\ &\geq 0 \quad (\text{since } n_l \geq 1) \end{aligned}$$

Second term:

$$\begin{aligned}
 \Delta^b &= \sum_{o \in \mathcal{O} \setminus c} \sum_{e \in \mathcal{E}} \left(-\frac{n^o}{N} \prod_{i \in e} n_i^o + \frac{\widehat{n}^o}{\widehat{N}} \prod_{i \in e} \widehat{n}_i^o \right) \\
 &= \sum_{o \in \mathcal{O} \setminus c} \sum_{e \in \mathcal{E}} \left(-\frac{n^o}{N} \prod_{i \in e} n_i^o + \frac{n^o}{N-1} \prod_{i \in e} n_i^o \right) \\
 &= \sum_{o \in \mathcal{O} \setminus c} \sum_{e \in \mathcal{E}} \frac{n^o}{N(N-1)} \prod_{i \in e} n_i^o \\
 &\geq 0 \quad (\text{since each term} \geq 0)
 \end{aligned}$$

Third term:

$$\begin{aligned}
 \Delta^c &= \sum_{e \in \mathcal{E}} \left(-\frac{n^c}{N} \prod_{i \in e} n_i^c + \frac{\widehat{n}^c}{\widehat{N}} \prod_{i \in e} \widehat{n}_i^c \right) \\
 &= -\frac{n^c}{N} \sum_{e \in \mathcal{E}} \prod_{i \in e} n_i^c + \frac{n^c - 1}{N-1} \left(\sum_{e \in \mathcal{E}} \left(\prod_{i \in e, i \neq l} n_i^c \right) (n_l^c - 1)^{|e_l|} \right) \\
 &= -\frac{n^c}{N} \sum_{e \in \mathcal{E}} \prod_{i \in e} n_i^c + \frac{n^c - 1}{N-1} \left(\sum_{e \in \mathcal{E}} \left(\prod_{i \in e, i \neq l} n_i^c \right) \left((n_l^c)^{|e_l|} - (n_l^c)^{|e_l|} + (n_l^c - 1)^{|e_l|} \right) \right) \\
 &= -\frac{n^c}{N} \sum_{e \in \mathcal{E}} \prod_{i \in e} n_i^c + \frac{n^c - 1}{N-1} \left(\sum_{e \in \mathcal{E}} \prod_{i \in e} n_i^c - \sum_{e \in \mathcal{E}_l} \left(\prod_{i \in e, i \neq l} n_i^c \right) \left((n_l^c)^{|e_l|} - (n_l^c - 1)^{|e_l|} \right) \right) \\
 &= -\frac{N - n^c}{N(N-1)} \sum_{e \in \mathcal{E}} \prod_{i \in e} n_i^c - \frac{n^c - 1}{N-1} \left(\sum_{e \in \mathcal{E}_l} \left(\prod_{i \in e, i \neq l} n_i^c \right) \left((n_l^c)^{|e_l|} - (n_l^c - 1)^{|e_l|} \right) \right) \\
 &\leq 0 \quad (\text{since each term} \leq 0)
 \end{aligned}$$

We also define:

$$\begin{aligned}
 \Delta^c &= \left(\frac{N - n^c}{N(N-1)} \right) \Delta_1^c + \left(\frac{n^c - 1}{N-1} \right) \Delta_2^c \\
 \Delta_1^c &= - \sum_{e \in \mathcal{E}} \prod_{i \in e} n_i^c \\
 \Delta_2^c &= - \left(\sum_{e \in \mathcal{E}_l} \left(\prod_{i \in e, i \neq l} n_i^c \right) \left((n_l^c)^{|e_l|} - (n_l^c - 1)^{|e_l|} \right) \right) \\
 \Delta^a &= \left(\frac{N(N - n^c)}{N(N-1)} + \frac{n^c - 1}{N-1} \right) \Delta^a \\
 &= \left(\frac{N - n^c}{N(N-1)} \right) \Delta_1^a + \left(\frac{n^c - 1}{N-1} \right) \Delta_2^a \\
 \Delta_1^a &= N \sum_{e \in \mathcal{E}_l} \left(\prod_{i \in e, i \neq l} n_i \right) \left(n_l^{|e_l|} - (n_l - 1)^{|e_l|} \right) \\
 \Delta_2^a &= \sum_{e \in \mathcal{E}_l} \left(\prod_{i \in e, i \neq l} n_i \right) \left(n_l^{|e_l|} - (n_l - 1)^{|e_l|} \right)
 \end{aligned}$$

The constants in front of the sum for Δ_1^a and Δ_2^c were from the equation, and Δ^a was split up to include the same constants. Now we will show that $\Delta_1^a + \Delta_1^c \geq 0$ and $\Delta_2^a + \Delta_2^c \geq 0$. We start with the latter:

$$\begin{aligned} \Delta_2^a + \Delta_2^c &= \sum_{e \in \mathcal{E}_l} \left[\left(\prod_{i \in e, i \neq l} n_i \right) (n_l^{|e_l|} - (n_l - 1)^{|e_l|}) - \left(\prod_{i \in e, i \neq l} n_i^c \right) ((n_l^c)^{|e_l|} - (n_l^c - 1)^{|e_l|}) \right] \\ &\geq \sum_{e \in \mathcal{E}_l} \left[\left(\prod_{i \in e, i \neq l} n_i \right) (n_l^{|e_l|} - (n_l - 1)^{|e_l|}) - \left(\prod_{i \in e, i \neq l} n_i \right) (n_l^{|e_l|} - (n_l - 1)^{|e_l|}) \right] \\ &= 0 \end{aligned} \quad (7)$$

Where (7) follows from $n_i \geq n_i^c \forall i$.

$$\begin{aligned} \Delta_1^a + \Delta_1^c &= N \sum_{e \in \mathcal{E}_l} \left(\prod_{i \in e, i \neq l} n_i \right) (n_l^{|e_l|} - (n_l - 1)^{|e_l|}) - \sum_{e \in \mathcal{E}} \prod_{i \in e} n_i^c \\ &\geq N \sum_{e \in \mathcal{E}_l} \left(\prod_{i \in e, i \neq l} n_i \right) (n_l^{|e_l|} - (n_l - 1)^{|e_l|}) - \sum_{e \in \mathcal{E}} \prod_{i \in e} n_i \end{aligned} \quad (8)$$

$$\geq N \sum_{e \in \mathcal{E}_l} \left(\prod_{i \in e, i \neq l} n_i \right) n_l^{|e_l|-1} - \sum_{e \in \mathcal{E}} \prod_{i \in e} n_i \quad (9)$$

$$\begin{aligned} &= (N - n_l) \sum_{e \in \mathcal{E}_l} \left(\prod_{i \in e, i \neq l} n_i \right) n_l^{|e_l|-1} + \sum_{e \in \mathcal{E}_l} \left(\prod_{i \in e, i \neq l} n_i \right) n_l^{|e_l|} - \sum_{e \in \mathcal{E}} \prod_{i \in e} n_i \\ &= (N - n_l) \sum_{e \in \mathcal{E}_l} \left(\prod_{i \in e, i \neq l} n_i \right) n_l^{|e_l|-1} - \sum_{e \in \overline{\mathcal{E}}_l} \prod_{i \in e} n_i \\ &\geq (N - n_l) \sum_{e \in \mathcal{E}_l} \prod_{i \in e, i \neq l} n_i - \sum_{e \in \overline{\mathcal{E}}_l} \prod_{i \in e} n_i \end{aligned} \quad (10)$$

Where (8) follows from $n_i \geq n_i^c \forall i$, (9) follows from $n_l^{|e_l|} - (n_l - 1)^{|e_l|} \geq n_l^{|e_l|} - n_l^{|e_l|-1}(n_l - 1) = n_l^{|e_l|-1}$, and (10) cancels edges in \mathcal{E}_l exactly, leaving only edges in $\overline{\mathcal{E}}_l$.

We again want to separate out terms that cancel. We define:

$$\begin{aligned} \mathcal{E}^{\widehat{k}} &= \{e : |e| = \widehat{k} \wedge \nexists j \text{ s.t. } \forall g \in e : g \subseteq r_j\} \\ \mathcal{E}^{\min} &= \{e : e \in \mathcal{E}, \nexists \widehat{e} \subset e : \widehat{e} \in \mathcal{E}^{k-1}\} \\ \overline{\mathcal{E}^{\min}} &= \mathcal{E} \setminus \mathcal{E}^{\min} \end{aligned}$$

We defined $\mathcal{E}^{\widehat{k}}$ as the hyperedges for any specified cardinality \widehat{k} . We call \mathcal{E}^{\min} the *minimal* hyperedges if k is the minimal cardinality at which these regions should be separated. Thus, these are the hyperedges where no subset of subregions $\{g_1 \dots g_{k-1}\} \subset e$ would have a separation hyperedge. All other hyperedges are called *non-minimal*. We also define $\mathcal{E}_l^{\min}, \overline{\mathcal{E}}_l^{\min}, \mathcal{E}_l^{\overline{\min}}, \overline{\mathcal{E}}_l^{\overline{\min}}$ as the minimal and non-minimal hyperedges of \mathcal{E}_l and $\overline{\mathcal{E}}_l$:

$$\begin{aligned} \mathcal{E}_l^{\min} &= \{e : e \in \mathcal{E}_l, \nexists \widehat{e} \subset e : \widehat{e} \in \mathcal{E}^{k-1}\} \\ \overline{\mathcal{E}}_l^{\min} &= \{e : e \in \overline{\mathcal{E}}_l, \nexists \widehat{e} \subset e : \widehat{e} \in \mathcal{E}^{k-1}\} \\ \mathcal{E}_l^{\overline{\min}} &= \mathcal{E}_l \setminus \mathcal{E}_l^{\min} \\ \overline{\mathcal{E}}_l^{\overline{\min}} &= \overline{\mathcal{E}}_l \setminus \overline{\mathcal{E}}_l^{\min} \end{aligned}$$

We also note that:

$$\begin{aligned} \sum_{e \in \overline{\mathcal{E}}_l^{\min}} \prod_{i \in e} n_i &\leq \sum_{g_j \in \mathcal{G} \setminus g_l} n_j \sum_{e \in \overline{\mathcal{E}}_l^{k-1}} \prod_{i \in e} n_i \\ &= (N - n_l) \sum_{e \in \overline{\mathcal{E}}_l^{k-1}} \prod_{i \in e} n_i \end{aligned}$$

For convenience, we define one additional set of hyperedges $\widehat{\mathcal{E}}_l$. These are hyperedges in \mathcal{E}_l such that no subset of $k-1$ elements which do not include k are in $\overline{\mathcal{E}}_l$.

$$\widehat{\mathcal{E}}_l = \{e : e \in \mathcal{E}_l \wedge \nexists e_{k-1} \subset e \text{ s.t. } e_{k-1} \in \overline{\mathcal{E}}_l^{k-1}\}$$

This enables us to split the set \mathcal{E}_l up into edges where $\overline{\mathcal{E}}_l^{k-1}$ are a subset, and $\widehat{\mathcal{E}}_l$. We note that since there is no region shared by all elements $e^{k-1} \in \overline{\mathcal{E}}_l^{k-1}$, then there will be no region shared by $e = e^{k-1} \cup g_l$. Thus, this will be an element of \mathcal{E}_l . This gives us:

$$\sum_{e \in \mathcal{E}_l} \prod_{i \in e, i \neq l} n_i = \sum_{e \in \overline{\mathcal{E}}_l^{k-1}} \prod_{i \in e} n_i + \sum_{e \in \widehat{\mathcal{E}}_l} \prod_{i \in e, i \neq l} n_i$$

Applying these:

$$\begin{aligned} \Delta_1^a + \Delta_1^c &\geq (N - n_l) \sum_{e \in \mathcal{E}_l} \prod_{i \in e, i \neq l} n_i - \sum_{e \in \overline{\mathcal{E}}_l} \prod_{i \in e} n_i \\ &= (N - n_l) \left(\sum_{e \in \overline{\mathcal{E}}_l^{k-1}} \prod_{i \in e} n_i + \sum_{e \in \widehat{\mathcal{E}}_l} \prod_{i \in e, i \neq l} n_i \right) - \sum_{e \in \overline{\mathcal{E}}_l^{\min}} \prod_{i \in e} n_i - \sum_{e \in \mathcal{E}_l^{\min}} \prod_{i \in e} n_i \\ &\geq (N - n_l) \left(\sum_{e \in \overline{\mathcal{E}}_l^{k-1}} \prod_{i \in e} n_i + \sum_{e \in \widehat{\mathcal{E}}_l} \prod_{i \in e, i \neq l} n_i \right) - (N - n_l) \sum_{e \in \overline{\mathcal{E}}_l^{k-1}} \prod_{i \in e} n_i - \sum_{e \in \overline{\mathcal{E}}_l^{\min}} \prod_{i \in e} n_i \\ &= (N - n_l) \left(\sum_{e \in \widehat{\mathcal{E}}_l} \prod_{i \in e, i \neq l} n_i \right) - \sum_{e \in \overline{\mathcal{E}}_l^{\min}} \prod_{i \in e} n_i \end{aligned}$$

At this point, we use the structure of our edge construction and definition of k to show this sum is ≥ 0 . We have a positive term, consisting of edges which include k , and a negative term, consisting of edges that do not include k . We will show that for every product in the negative term, there is a corresponding product in the positive term.

To do so, we show that for any $e \in \overline{\mathcal{E}}_l^{\min}$, there is at least one corresponding $e' \in \widehat{\mathcal{E}}_l$ to cancel the terms out. More concretely:

Lemma 3.1.1. *Let $e \in \overline{\mathcal{E}}_l^{\min}$. There exists some $e^{k-1} \subset e$, $|e^{k-1}| = k-1$ such that $e' = (e^{k-1} \cup g_l) \in \widehat{\mathcal{E}}_l$.*

Proof. Recall that e is a multiset of subregions. It is straightforward to see that because e is minimal, there can be no repeated elements in the multiset - and thus it is equivalent to a set. Define this set as $e = \{\widehat{g}_1 \dots \widehat{g}_k\}$. Define each distinct subset which does not include \widehat{g}_i as $e_i = e \setminus \widehat{g}_i$, $1 \leq i \leq k$. By our definition of minimal hyperedges $\overline{\mathcal{E}}_l^{\min}$, we know that $\forall e_i, \exists r_i : e_i \subseteq r_i$, which implies that $e_i \not\subseteq \overline{\mathcal{E}}_l^{k-1}$. Note that each r_i must be distinct. If $r_i = r_j$, for any i, j , this would imply $(e_i \cup e_j) = e \subseteq r_i$. But since there exists a separating hyperedge $e, \exists r : e \subseteq r$. This implies $\widehat{g}_i \not\subseteq r_i$. Combining this with our definition of $\widehat{\mathcal{E}}_l$, if $\exists r : (e_i \cup g_l) \subseteq r \Rightarrow (e_i \cup g_l) \in \widehat{\mathcal{E}}_l$. To prove this lemma, we will show that this region cannot exist for all e_i .

If $g_l \not\subseteq r_i \Rightarrow e_i \cup g_l \not\subseteq r_i$. For the sake of contradiction, suppose $g_l \subseteq r_i \forall i$. Let $R = \{r_1 \dots r_k\}$. For this to be true, it must be that: ① $\forall h \in g_l, h \in R$ ② $\forall r_i \in R, \forall \widehat{h} \in \widehat{g}_i : \widehat{h} \notin r_i, \widehat{h} \in R \setminus r_i$ where $|R| = k$. However, by definition of k this cannot be true: the largest such R where this holds $|R| = k-1$. Thus, we have a contradiction, and have shown such a set of regions $\{r_1 \dots r_k\} = R : g_l \subseteq r_i \forall r_i$ cannot exist. Therefore, $\exists e_i : (e_i \cup g_l) \in \widehat{\mathcal{E}}_l$. \square

In order to apply Lemma 3.1.1, we split every $e \in \overline{\mathcal{E}_l^{\min}}$ it up into e^{k-1} and \bar{g} , where e^{k-1} is the subset of e such that $(e^{k-1} \cup g_l) \in \widehat{\mathcal{E}}_l$, and $\bar{g} = e \setminus e^{k-1}$. Let \bar{n} be the number of particles in subregion \bar{g} , which we will use in ((11)):

$$\begin{aligned} \Delta_1^a + \Delta_1^c &\geq (N - n_l) \left(\sum_{e \in \widehat{\mathcal{E}}_l} \prod_{i \in e, i \neq l} n_i \right) - \sum_{e \in \overline{\mathcal{E}_l^{\min}}} \prod_{i \in e} n_i \\ &= (N - n_l) \left(\sum_{e \in \widehat{\mathcal{E}}_l} \prod_{i \in e, i \neq l} n_i \right) - \sum_{e \in \overline{\mathcal{E}_l^{\min}}} \bar{n} \prod_{i \in e^{k-1}} n_i \end{aligned} \quad (11)$$

$$\begin{aligned} &\geq (N - n_l) \left(\sum_{e \in \widehat{\mathcal{E}}_l} \prod_{i \in e, i \neq l} n_i \right) - \sum_{g_j \in \mathcal{G} \setminus g_l} n_j \left(\sum_{e \in \widehat{\mathcal{E}}_l} \prod_{i \in e, i \neq l} n_i \right) \\ &= (N - n_l) \left(\sum_{e \in \widehat{\mathcal{E}}_l} \prod_{i \in e, i \neq l} n_i \right) - (N - n_l) \left(\sum_{e \in \widehat{\mathcal{E}}_l} \prod_{i \in e, i \neq l} n_i \right) \\ &= 0 \end{aligned} \quad (12)$$

Where ((12)) applies Lemma 3.1.1.

At this point, we have shown that $\Delta = \Delta^a + \Delta^b + \Delta^c \geq 0$, since $\Delta^b \geq 0$ and $\Delta^a + \Delta^c \geq 0$, which is what we needed to show. □

It is not hard to see that for any $\mathcal{S} \subseteq \widehat{\mathcal{S}} \subseteq \mathcal{T} \times \mathcal{O}$, we could show that $\Delta(t|\mathcal{S}) \geq \Delta(t|\widehat{\mathcal{S}}_1) \geq \Delta(t|\widehat{\mathcal{S}}_2) \cdots \geq \Delta(t|\widehat{\mathcal{S}})$. In other words, we can always find a sequence of removing one hypothesis at a time to get from \mathcal{S} to $\widehat{\mathcal{S}}$ when $\mathcal{S} \subseteq \widehat{\mathcal{S}} \subseteq \mathcal{T} \times \mathcal{O}$. □

7.4 Theorem 3: Greedy Performance Bound

We would like to apply Theorem 5.8 of Golovin and Krause (2011). We have already shown adaptive submodularity and strong adaptive monotonicity in Sec. 7.3. The theorem also requires that instances are *self-certifying*, which means that when the policy knows it has obtained the maximum possible objective value immediately upon doing so. See Golovin and Krause (2011) for details. As our objective is equivalent for all remaining hypotheses in $\mathcal{V}(\mathcal{S})$, our function f_{HEC} is self-certifying. □

The performance bound now follows directly from Theorem 5.8 of Golovin and Krause (2011). To apply the theorem, we needed to define two constants: a bound on the maximum value of $f_{\text{HEC}}(\mathcal{S})$, $Q = 1$, and the minimum our objective function can change by, which corresponds to removing one hyperedge, $\eta = p_{\min}^k$. Plugging those into Theorem 5.8 of Golovin and Krause (2011) gives $\mathcal{C}(\pi_{\text{HEC}}) \leq (k \ln(1/p_{\min}) + 1)\mathcal{C}(\pi^*)$. □