

PERFORMANCE GUARANTEES FOR INFORMATION THEORETIC SENSOR RESOURCE MANAGEMENT

Jason L. Williams,^{1,2} John W. Fisher III,² Alan S. Willsky^{1,2}

¹Laboratory for Information and Decision Systems and

²Computer Science and Artificial Intelligence Laboratory

Massachusetts Institute of Technology, Cambridge MA 02139

jlwilliams@alum.mit.edu fisher@csail.mit.edu willsky@mit.edu

ABSTRACT

Many estimation problems involve sensors which can be actively controlled to alter the information received and utilized in the underlying inference task. In this paper, we discuss performance guarantees for heuristic algorithms for adaptive sensor control in sequential estimation problems, where the inference criterion is mutual information. We also demonstrate the performance of our tighter online computable performance guarantees through computational simulations. The guarantees may be applied to other estimation criteria including the Cramér-Rao bound.

Index Terms— Sequential decision procedures, sequential estimation, tracking

1. INTRODUCTION

Active sensing is motivated by modern sensors which can be controlled to observe different aspects of an underlying probabilistic process. For example, if we use cameras to track people in buildings, we can steer the camera to focus or zoom on different people or places; in a sensor network, we can choose to activate and deactivate different nodes and different sensing modalities within a particular node; or in a medical diagnosis problem we can choose which tests to administer to a patient. In each of these cases, our control choices impact the information that we receive in our observation, and thus the performance achieved in the underlying inference task.

A commonly used performance objective in active sensing is mutual information (MI) (*e.g.*, [1]). Denoting the quantity that we are aiming to infer as X and the observation resulting from control choice u as z^u , the MI between X and z^u is defined as the expected reduction in the entropy produced by the observation [2], *i.e.*, $I(X; z^u) = H(X) - H(X|z^u) = H(z^u) - H(z^u|X)$.¹ Since $H(X)$ is independent of the control choice u , choosing u to maximize $I(X; z^u)$ is equivalent

to minimizing the uncertainty in X as measured by the conditional entropy $H(X|z^u)$.

In different problems, the collection of subsets of observations from which one may choose can have a dramatically different structure. One common structure involves selection of any K -element subset of a set of observations, *e.g.*, the subset of sensors to activate in a sensor network application. Another structure is one in which there is a single sensor which can operate in one mode at each time increment; the resulting selection structure is one in which we may choose a single element from each of a series of observation sets, each of which corresponds to a different time instant.

Recent work [3] has applied results from [4] to establish that, when the selection structure is such that any subset of observations with cardinality $\leq K$ may be chosen, the greedy heuristic (which at each stage chooses the observation which maximizes the MI with X conditioned on the already selected observations) achieves a total MI of within a constant multiple $(1 - 1/e) \approx 0.632$ of the optimal subset of observations. Our analysis extends this to the larger class of problems involving sequential processes, providing the surprising result that in sequential problems, under quite general assumptions one may select the control for the current time instant neglecting future observation opportunities, and still have performance $\geq 0.5 \times$ the optimal. Furthermore, the online computable bounds demonstrated in Section 2.3 can be significantly stronger in certain circumstances. In the interest of space, most proofs are omitted; details can be found in [5].

The guarantees we develop are based upon submodularity, the same property exploited in [3, 4, 6]. Submodularity captures the notion that as we select more observations, the value of the remaining unselected observations decreases, *i.e.*, the notion of diminishing returns.

Definition 1. A set function f is submodular if $f(\mathcal{C} \cup \mathcal{A}) - f(\mathcal{A}) \geq f(\mathcal{C} \cup \mathcal{B}) - f(\mathcal{B}) \forall \mathcal{B} \supseteq \mathcal{A}$.

It was established in [3] that, assuming that observations are independent conditioned on the quantity to be estimated, MI is a submodular function of the observation selection set. The simple result that we will utilize from submodularity is that $I(x; z^{\mathcal{C}}|z^{\mathcal{A}}) \geq I(x; z^{\mathcal{C}}|z^{\mathcal{B}}) \forall \mathcal{B} \supseteq \mathcal{A}$.

This work was supported by MIT Lincoln Laboratory through ACC PO#3019934.

¹Note that when we condition on a random variable (such as a yet unrealized observation) the conditional entropy involves an expectation over the distribution of that random variable.

2. A SIMPLE PERFORMANCE GUARANTEE

To commence, consider a simple problem involving two time steps, where at each step we must choose a single observation from a set (e.g., in which mode to operate a sensor). The goal is to maximize the information obtained about an underlying quantity X . Let $\{o_1, o_2\}$ denote the optimal choice for the two stages, i.e., that which maximizes $I(X; z_1^{u_1}, z_2^{u_2})$ over possible choices for $\{u_1, u_2\}$. Let $\{g_1, g_2\}$ denote the choice made by the greedy heuristic, where $g_1 = \arg \max_{u_1} I(X; z_1^{u_1})$ and $g_2 = \arg \max_{u_2} I(X; z_2^{u_2} | z_1^{g_1})$ (where conditioning is on the random variable $z_1^{g_1}$, not on the resulting observation value). Then the following analysis establishes a performance guarantee for the greedy algorithm:

$$\begin{aligned}
 I(X; z_1^{o_1}, z_2^{o_2}) &\stackrel{(a)}{\leq} I(X; z_1^{g_1}, z_2^{g_2}, z_1^{o_1}, z_2^{o_2}) \\
 &\stackrel{(b)}{=} I(X; z_1^{g_1}) + I(X; z_2^{g_2} | z_1^{g_1}) \\
 &\quad + I(X; z_1^{o_1} | z_1^{g_1}, z_2^{g_2}) \\
 &\quad + I(X; z_2^{o_2} | z_1^{g_1}, z_2^{g_2}, z_1^{o_1}) \\
 &\stackrel{(c)}{\leq} I(X; z_1^{g_1}) + I(X; z_2^{g_2} | z_1^{g_1}) \\
 &\quad + I(X; z_1^{o_1}) + I(X; z_2^{o_2} | z_1^{g_1}) \\
 &\stackrel{(d)}{\leq} 2I(X; z_1^{g_1}) + 2I(X; z_2^{g_2} | z_1^{g_1}) \\
 &\stackrel{(e)}{=} 2I(X; z_1^{g_1}, z_2^{g_2}) \tag{1}
 \end{aligned}$$

where (a) results from the nondecreasing property of MI, (b) is an application of the MI chain rule, (c) results from submodularity (assuming that all observations are independent conditioned on X), (d) from the definition of the greedy heuristic, and (e) from a reverse application of the chain rule. Thus the optimal performance can be no more than twice that of the greedy heuristic, or, conversely, the performance of the greedy heuristic is at least half that of the optimal.²

Theorem 1 presents this result in its most general form; the proof directly follows the above steps. The following assumption establishes the basic structure: we have N sets of observations, and we can select a specified number of observations from each set in an arbitrary order.

Assumption 1. *There are N sets of observations, $\{\{z_1^1, \dots, z_1^{n_1}\}, \{z_2^1, \dots, z_2^{n_2}\}, \dots, \{z_N^1, \dots, z_N^{n_N}\}\}$, which are mutually independent conditioned on the quantity to be estimated (X). Any k_i observations can be chosen out of the i -th set ($\{z_i^1, \dots, z_i^{n_i}\}$). The observation sets are visited using the greedy algorithm in an order specified by a sequence $(w_1, \dots, w_M)$ where $w_i \in \{1, \dots, N\} \forall i$ (i.e., in the i -th stage we select a previously unselected observation out of the w_i -th set).*

²Note that this is considering only open loop control; we will discuss closed loop control in Section 3.

The abstraction of the observation set sequence $(w_1, \dots, w_M)$ allows us to visit observation sets more than once (allowing us to select multiple observations from each set) and in any order. The greedy heuristic operating on this structure is defined below, followed by the general form of the guarantee.

Definition 2. *The greedy heuristic operates according to the following rule:*

$$g_j = \arg \max_{u \in \{1, \dots, n_{w_j}\}} I(X; z_{w_j}^u | z_{w_1}^{g_1}, \dots, z_{w_{j-1}}^{g_{j-1}})$$

Theorem 1. *Under Assumption 1, the greedy heuristic in Definition 2 has performance guaranteed by the following expression:*

$$I(X; z_{w_1}^{o_1}, \dots, z_{w_M}^{o_M}) \leq 2I(X; z_{w_1}^{g_1}, \dots, z_{w_M}^{g_M})$$

where $\{z_{w_1}^{o_1}, \dots, z_{w_M}^{o_M}\}$ is the optimal set of observations, i.e., the one which maximizes $I(X; z_{w_1}^{u_1}, \dots, z_{w_M}^{u_M})$ over the possible choices for $\{u_1, \dots, u_M\}$.

The proof of the theorem can be found in [5], where it is also shown that the bound is tight.

2.1. Comparison to matroid guarantee

The prior work using matroids [6] provides another algorithm with the same guarantee for problems of this structure. However, to achieve the guarantee on matroids it is necessary to consider every observation at every stage of the problem. Computationally, it is far more desirable to be able to proceed in a dynamic system by selecting observations at time k considering only the observations available at that time, disregarding future time steps (indeed, countless previous works, such as [1] do just that). The freedom of choice of the order in which we visit observation sets in Theorem 1 extends the performance guarantee to this commonly used sequential selection structure.

2.2. Online version of guarantee

Modifying step (c) of Eq. (1) (or the corresponding step of the proof of Theorem 1), we can also obtain an online performance guarantee, which will often be substantially tighter in practice (as demonstrated in Section 2.3). The online bound can be used to calculate an upper bound for the optimal reward starting from any sequence of observation choices, not just the choice made by the greedy heuristic in Definition 2, $(g_1, \dots, g_M)$. The online bound will tend to be tight in cases where the amount of information remaining after choosing the set of observations is small.

Theorem 2. *Under the same assumptions as Theorem 1, for each $i \in \{1, \dots, N\}$ define $\bar{k}_i = \min\{k_i, n_i - k_i\}$, and for each $j \in \{1, \dots, \bar{k}_i\}$ define*

$$\bar{g}_i^j = \arg \max_{u \in \{1, \dots, n_i\} - \{\bar{g}_i^l | l < j\}} I(X; z_i^u | z_{w_1}^{g_1}, \dots, z_{w_M}^{g_M}) \tag{2}$$

Then the following two performance guarantees, which are computable online, apply:

$$\begin{aligned} I(X; z_{w_1}^{o_1}, \dots, z_{w_M}^{o_M}) \\ \leq I(X; z_{w_1}^{g_1}, \dots, z_{w_M}^{g_M}) \\ + \sum_{i=1}^N \sum_{j=1}^{\bar{k}_i} I(X; z_i^{\bar{g}_i^j} | z_{w_1}^{g_1}, \dots, z_{w_M}^{g_M}) \end{aligned} \quad (3)$$

$$\begin{aligned} \leq I(X; z_{w_1}^{g_1}, \dots, z_{w_M}^{g_M}) \\ + \sum_{i=1}^N \bar{k}_i I(X; z_i^{\bar{g}_i^1} | z_{w_1}^{g_1}, \dots, z_{w_M}^{g_M}) \end{aligned} \quad (4)$$

2.3. Example of online guarantee

Suppose that we are using a surface vehicle travelling at a constant velocity along a fixed path (as illustrated in Fig. 1(a)) to map the depth of the ocean floor in a particular region. Assume that, at any position on the path (such as the points denoted by ‘ $\triangle$ ’), we may steer our sensor to measure the depth of any point within a given region around the current position (as depicted by the dotted ellipses), and that we receive a linear measurement of the depth corrupted by Gaussian noise with variance R . Suppose that we model the depth of the ocean floor as a Gauss-Markov random field with a 500×100 thin membrane grid model where neighboring node attractions are uniformly equal to q . One cycle of the vehicle path takes 300 time steps to complete.


Defining the state X to be the vector containing one element for each cell in the 500×100 grid, the problem can be seen to fit into the structure of Assumption 1 (with $w_i = i$ and $k_i = 1 \forall i$). The selection algorithm simply selects at each stage the most informative observation conditioned on the observations previously chosen. A single observation of the cell directly beneath the sensing platform is used as initialization to obtain a full rank information matrix.

Fig. 1(b) shows the accrual of reward over time as well as the bound on the optimal sequence obtained using Theorem 2 for each time step when $q = 100$ and $R = 1/40$, while Fig. 1(c) shows the ratio between the achieved performance and the optimal sequence bound over time. The graph indicates that the greedy heuristic achieves at least $0.8 \times$ the optimal reward. The tightness of the online bound depends on particular model characteristics: if $q = R = 1$, then the guarantee ratio is much closer to the value of the offline bound (*i.e.*, 0.5).


2.4. Exploiting diffusiveness

In problems such as object tracking, the kinematic quantities of interest evolve according to a diffusive process, in which correlation between states at different time instants reduces as the time difference increases. Intuitively, one would expect that a greedy algorithm would be closer to optimal in

(a) Region boundary and vehicle path



(b) Reward of greedy heuristic and bound on optimal



(c) Factor of optimality from online guarantee




Fig. 1. (a) shows region boundary and vehicle path (counter-clockwise, starting from the left end of the lower straight segment). When the vehicle is located at a ‘ $\triangle$ ’ mark, any one grid element with center inside the surrounding dotted ellipse may be measured. (b) graphs reward accrued by the greedy heuristic after different periods of time, and the bound on the optimal sequence for the same time period. (c) shows the ratio of these two curves, providing the factor of optimality guaranteed by the bound.

situations in which the diffusion strength is high. In [5], we explore an extension of Theorems 1 and 2 which exploits diffusiveness in order to obtain tighter guarantees (in the offline form of Theorem 1, as well as in the online computable form of Theorem 2). In the interest of space we omit details from this presentation.

3. CLOSED LOOP CONTROL

The analysis in Sections 2 and 2.4 concentrates on an open loop control structure, *i.e.*, it assumes that all observation choices are made before any observation values are received. Greedy heuristics are often applied in a closed loop setting, in which an observation is chosen, and then its value is received before the next choice is made.

The performance guarantee of Theorem 1 applies to the *expected* performance of the greedy heuristic operating in a closed loop fashion, *e.g.*, in expectation the closed loop greedy heuristic achieves at least half the reward of the optimal *open* loop selection. We refer the reader to [5] for a proof of this result. We emphasize that this performance guarantee is for *expected* performance: it does not provide a guarantee for the change in entropy of every sample path. An online bound cannot be obtained on the basis of a single realization, although online bounds can be calculated through Monte Carlo simulation (to approximate the expectation).

There is no guarantee relating the performance of the closed loop greedy heuristic to the optimal closed loop controller, as the following example illustrates. One exception to this is linear Gaussian models, where closed loop policies can perform no better than open loop sequences, so that the open loop guarantee extends to closed loop performance.

Example 1. Consider the following two-stage problem, where $X = [a, b, c]^T$, with $a \in \{1, \dots, N\}$, $b \in \{1, \dots, N+1\}$, and $c \in \{1, \dots, M\}$. The prior distribution of each of these is uniform and independent. In the first stage, we may measure $z_1^1 = a$ for reward $\log N$, or $z_1^2 = b$ for reward $\log(N+1)$. In the second stage, we may choose z_2^i , $i \in \{1, \dots, N\}$, where

$$z_2^i = \begin{cases} c, & i = a \\ d, & \text{otherwise} \end{cases}$$

where d is independent of X , and is uniformly distributed on $\{1, \dots, M\}$. The greedy algorithm in the first stage selects the observation $z_1^2 = b$, as it yields a higher reward ($\log(N+1)$) than $z_1^1 = a$ ($\log N$). At the second stage, all options have the same reward, $\frac{1}{N} \log M$, so we choose one arbitrarily for a total reward of $\log(N+1) + \frac{1}{N} \log M$. The optimal algorithm in the first stage selects the observation $z_1^1 = a$ for reward $\log N$, followed by the observation z_2^a for reward $\log M$, for total reward $\log N + \log M$. The ratio of the greedy reward to the optimal reward is

$$\frac{\log(N+1) + \frac{1}{N} \log M}{\log N + \log M} \rightarrow \frac{1}{N}, \quad M \rightarrow \infty$$

Hence, by choosing N and M to be large, we can obtain an arbitrarily small ratio between the greedy closed-loop reward and the optimal closed-loop reward.

4. GUARANTEES ON THE CRAMÉR-RAO BOUND

While the preceding discussion has focused exclusively on mutual information, the results are applicable to a larger class of objectives, including the log determinant of Fisher information, from which a guarantee may be derived for the posterior Cramér-Rao bound (PCRB), defined in [7]. If we denote

by C_X^A the PCRB matrix (*i.e.*, the inverse of the Fisher information) for estimating X using the set of observations z^A , the result of Theorem 1 can be used to show that:

$$|C_X^G| \leq |C_X^\emptyset| \sqrt{\frac{|C_X^\emptyset|}{|C_X^\emptyset|}}$$

where C_X^G is the PCRB matrix for X using the observations chosen by the greedy heuristic, and $C_X^\emptyset$ is the PCRB matrix using the optimal choice of observations (where the objective is the determinant of the Fisher information). Again, we defer proof of the result to [5]. The ratio $|C_X^G|/|C_X^\emptyset|$ is the fractional reduction of uncertainty (measured through covariance determinant) which is gained through using the selected observations rather than the prior information alone. Thus the result provides a guarantee on how much of the optimal reduction you lose by using the greedy heuristic. The determinant of the error covariance of any estimator of X using the data z^G is lower bounded by $|C_X^G|$.

5. CONCLUSION

The performance guarantees presented in this paper provide theoretical basis for simple heuristic algorithms that are widely used in practice. The guarantees apply to both open loop and closed loop operation, and are naturally tighter for diffusive processes. Substantially stronger online guarantees can be obtained for specific problems through computation of additional quantities after the greedy selection has been completed.

6. REFERENCES

- [1] K.J. Hintz and E.S. McVey, "Multi-process constrained estimation," *Systems, Man and Cybernetics, IEEE Transactions on*, vol. 21, no. 1, pp. 237–244, 1991.
- [2] Thomas M. Cover and Joy A. Thomas, *Elements of Information Theory*, John Wiley and Sons, New York, NY, 1991.
- [3] Andreas Krause and Carlos Guestrin, "Near-optimal nonmyopic value of information in graphical models," in *UAI 2005*, July 2005.
- [4] G.L. Nemhauser, L.A. Wolsey, and M.L. Fisher, "An analysis of approximations for maximizing submodular set functions—I," *Mathematical Programming*, vol. 14, no. 1, pp. 265–294, December 1978.
- [5] Jason L. Williams, *Information Theoretic Sensor Management*, Ph.D. thesis, Massachusetts Institute of Technology, February 2007.
- [6] G.L. Nemhauser, L.A. Wolsey, and M.L. Fisher, "An analysis of approximations for maximizing submodular set functions—II," in *Polyhedral combinatorics*, M.L. Balinski and A.J. Hoffman, Eds., vol. 8 of *Mathematical programming study*, pp. 73–87. Elsevier, 1978.
- [7] Harry L. Van Trees, *Detection, Estimation, and Modulation Theory*, Wiley-Interscience, 2001.