

THEORETICAL GUARANTEES ON PENALIZED INFORMATION GATHERING

Georgios Papachristoudis and John W. Fisher III

Massachusetts Institute of Technology

ABSTRACT

Optimal measurement selection for inference is combinatorially complex and intractable for large scale problems. Under mild technical conditions, it has been proven that greedy heuristics combined with conditional mutual information rewards achieve performance within a factor of the optimal. Here we provide conditions under which cost-penalized mutual information may achieve similar guarantees. Specifically, if the cost of a measurement is proportional to the information it conveys, the bounds proven in [4] and [10] still apply.

Index Terms— information measures, sensor selection

1. INTRODUCTION

Information gathering subject to resource constraints imposes significant complexities on modern sensing and actuation systems. Information-driven methods seek to maximize information extraction while limiting resource expenditures via active control of the measurement process. Recent sensor signal processing methods (e.g., [1, 9, 11]) consider mutual information as the reward embedded in a dynamic sensing algorithm. The problem of choosing an optimal subset of measurements is formulated as a combinatorial optimization problem which becomes intractable as the number of measurements grows. Utilizing results from Nemhauser *et al* [7], [4, 3, 10] provide theoretical performance bounds that guarantee greedy measurement selection to be within a computable factor of the optimal selection choice. The bounds depend critically upon whether the information rewards are *submodular*. Whereas the previous results consider selection constraints on pure information rewards, here we consider explicit costs/penalties to the reward function. We show a class of penalized rewards for which submodularity is preserved and thus the guarantees of [3, 4, 10] are applicable.

1.1. Problem Description

Consider inference of n latent variables $\mathbf{X} = \{X_1, \dots, X_n\}$ from measurements $\mathbf{Z} = \{Z_1, \dots, Z_n\}$. Each Z_t is a vector comprising of N_t measurements. The above can be easily

extended to the multidimensional case, where each element of Z_t would represent a multidimensional measurement of X_t . We will denote the indices of these measurements by $\mathcal{V}^t = \{1, \dots, N_t\}, \forall t \in \{1, \dots, n\}$ and refer to them as *measurement sets* herein. In addition, we denote by $\mathcal{V} = \cup_{t=1}^n \mathcal{V}^t$.

A common assumption, and necessary to the guarantees herein, is that measurements are statistically independent *conditioned* on \mathbf{X} (i.e., $Z_i \perp\!\!\!\perp Z_j | \mathbf{X}$). Trivially, the results extend to the more general case, where the number of latent variables and measurement sets is not necessarily the same, so long as *conditional* independence of measurements is satisfied. We describe performance bounds for greedy selection of measurements in both the batch and sequential settings. The *batch* setting treats \mathbf{X} as a monolithic hidden variable and \mathbf{Z} as a uniform collection while the *sequential* setting exploits the dependency structure of $\{X_1, \dots, X_n\}$ and the selection structure of $\{Z_1, \dots, Z_n\}$.

Consequently, in the *batch* setting there are $N = \sum_{t=1}^n N_t$ measurements, that are simultaneously available and the goal is to find the best subset up to size K , where $K \leq N$ under an information reward. In the *sequential* setting, there are N_t measurements for each hidden state X_t . Additionally, let $\{w_1, \dots, w_M\}$ denote the order in which the \mathcal{V}^t are visited during greedy selection. Clearly, $w_l \in \{1, \dots, n\}, \forall l \in \{1, \dots, M\}$. In addition, $\sum_{l=1}^M \delta(w_l, t) \leq k_t, \forall t$ and $\sum_{t=1}^n k_t = K$, where $\delta(\cdot, \cdot)$ is the Kronecker delta function. The last constraint implies that at most k_t measurements are chosen from the measurement set \mathcal{V}^t . Remarkably, in the sequential setting the greedy heuristic may be implemented without revealing the entire problem structure (e.g., if the walk corresponds to time) while still yielding performance bounds. By contrast, the entire problem structure must be available during greedy selection in the batch setting.

Furthermore, we consider a reward function $f : 2^{\mathcal{V}} \rightarrow \mathbb{R}$, that captures the value of sensing actions, and a cost function $c : 2^{\mathcal{V}} \rightarrow \mathbb{R}_+$ both defined on all subsets of measurements, \mathcal{V} . The evaluation of a single element j , $c(j)$ denotes the cost of the corresponding measurement. We will assume that the reward function is submodular.

Definition 1. Submodular Function. Given a finite set \mathcal{V} , a real-valued function f on the set of subsets of \mathcal{V} is submodular if

$$f(\mathcal{A}) + f(\mathcal{B}) \geq f(\mathcal{A} \cup \mathcal{B}) + f(\mathcal{A} \cap \mathcal{B}), \forall \mathcal{A}, \mathcal{B} \subseteq \mathcal{V}.$$

GP was partially supported by the Army Research Office MURI program #W911NF-11-1-0391. JF was partially supported by DARPA award #FA8650-11-1-7154

Additionally, we define the *set increment function* as the incremental value of adding element j to a set \mathcal{S} as:

$$\rho_{\mathcal{S}}(j) \triangleq f(\mathcal{S} \cup \{j\}) - f(\mathcal{S}) .$$

It trivially holds that: $\rho_{\emptyset}(j) = f(j)$. Equivalently, the set increment function of a submodular function satisfies:

$$\rho_{\mathcal{A}}(j) \geq \rho_{\mathcal{B}}(j) , \forall \mathcal{A} \subseteq \mathcal{B} \subseteq \mathcal{V} \text{ and } j \notin \mathcal{B} .$$

Submodularity, a discrete analogue to convexity [2], captures the notion of “diminishing returns”. Specifically, when the set increment function is the conditional mutual information, the *expected* information contained in a prospective measurement is higher when conditioned on the set of measurements \mathcal{A} than set \mathcal{B} , $\forall \mathcal{A} \subseteq \mathcal{B}$: $\rho_{\mathcal{A}}(j) = I(\mathbf{X}; Z^j | Z^{\mathcal{A}}) \geq I(\mathbf{X}; Z^j | Z^{\mathcal{B}}) = \rho_{\mathcal{B}}(j)$, where $Z^{\mathcal{S}}$ is defined as $Z^{\mathcal{S}} \triangleq \{Z^j | j \in \mathcal{S}\}$ for some $\mathcal{S} \subseteq \mathcal{V}$ [5]. Costs are assumed to be non-negative and additive

$$c(\mathcal{S}) = \sum_{j \in \mathcal{S}} c(j) . \quad (1)$$

Under the assumption of conditionally independent measurements, conditional mutual information is a *monotone* submodular function [4]. Such functions have the property that $f(\mathcal{A}) \leq f(\mathcal{B})$, $\forall \mathcal{A} \subseteq \mathcal{B}$, or equivalently $\rho_{\mathcal{S}}(j) \geq 0$, $\forall j \in \mathcal{V}$, $\mathcal{S} \subseteq \mathcal{V}$.

We define the greedy heuristic in the batch setting as

$$g_j = \arg \max_{u \in \{1, \dots, n\} \setminus \mathcal{G}^{j-1}} I(\mathbf{X}; Z^u | Z^{g_1}, \dots, Z^{g_{j-1}}) .$$

That is, choose the measurement with highest reward at step j conditioned on previous measurements. We denote the greedy solution at step $(j-1)^{\text{th}}$ as $\mathcal{G}^{j-1} \triangleq g_{1:j-1}$. Via [7], the greedy solution over monotone submodular functions is guaranteed to be at least $(1 - 1/e) \approx 0.63$ of the optimal reward.

Greedy selection in the sequential setting is considered in [10]. The essential difference is that sets of measurements are visited in sequence with selection restricted to the current set. In this case, greedy selection is defined as

$$g_j = \arg \max_{u \in \{1, \dots, n_{w_j}\} \setminus \mathcal{G}^{j-1}} I(\mathbf{X}; Z_w^u | Z_{w_1}^{g_1}, \dots, Z_{w_{j-1}}^{g_{j-1}}) .$$

While the proof method of [7] does not extend to the sequential setting, [8, 10] showed that the reward of the greedy heuristic is no worse than 1/2 of the optimal for any visiting sequence when the reward function is monotone and submodular. We next discuss conditions under which penalized information rewards may achieve similar guarantees.

2. METHOD

A natural approach is to consider penalized information rewards where the expected information gain is discounted by a cost function multiplied by a conversion factor. A difficulty

arises in that such functions may not be submodular and furthermore may not be monotone, rendering the previous guarantees inapplicable. Consistent with the above, we define the batch greedy heuristic for non-monotone submodular functions as follows:

$$g_j = \arg \max_{u \in \{1, \dots, n\} \setminus \mathcal{G}^{j-1}} \rho_{\mathcal{G}^{j-1}}(u) , \quad (2)$$

and similarly the sequential greedy heuristic as:

$$g_j = \arg \max_{u \in \{1, \dots, n_{w_j}\} \setminus \mathcal{G}^{j-1}} \rho_{\mathcal{G}^{j-1}}(u) .$$

Consider the subset of size less or equal to K that solves the following optimization:

$$\mathcal{O} = \arg \max_{\mathcal{S} \subseteq \mathcal{V}, |\mathcal{S}| \leq K} f(\mathcal{S}) \quad (3)$$

We assume that $f(\emptyset) = 0$ and that the greedy heuristic given in (2) terminates in L steps. Following Nemhauser *et al* [7] and assuming a worst-case negative reward, $\rho_{\mathcal{S}}(j) \geq -\theta$, $\forall j \in \mathcal{V}$, $\mathcal{S} \subseteq \mathcal{V}$, the following bound holds:

$$f(\mathcal{G}) \geq \left[1 - \frac{L}{K} \left(1 - \frac{1}{K} \right)^L \right] f(\mathcal{O}) - L\theta \left(1 - \frac{1}{K} \right)^L . \quad (4)$$

As K grows, this bound is of less utility.

As stated, the results of [10] apply to the sequential setting when the reward is submodular and monotone. Unfortunately, the latter property is not satisfied here. However, one can introduce k_t *auxiliary* measurements for each measurement set \mathcal{V}^t , $\forall t \in \{1, \dots, n\}$ that have no information value and incur no cost. Such measurements are equivalent to choosing no measurements at a given step. Introduction of these measurements guarantees that greedy selection becomes nonnegative and that the greedy solution terminates in exactly K steps. Let us further assume that the optimal solution \mathcal{O} has size K^* , where $K^* \leq K$. We can introduce $K - K^*$ auxiliary variables, and thus obtain an optimal solution of size K with the same reward as the one of size K^* .

Then, defining the function $f'(\mathcal{S}) = f(\mathcal{S}) + |\mathcal{S}|\theta$, it is easily shown that $f'(\cdot)$ is both submodular and monotone. Thus the known bound applies:

$$f'(\mathcal{O}) \leq 2f'(\mathcal{G}) \Leftrightarrow f(\mathcal{O}) \leq 2f(\mathcal{G}) + K\theta .$$

where \mathcal{O} and \mathcal{G} are the optimal and greedy selections under the original $f(\cdot)$.

2.1. Penalized Mutual Information

Consider the following reward function

$$f(\mathcal{S}) = \max \{ I(\mathbf{X}; Z^{\mathcal{S}}) - \lambda c(\mathcal{S}), 0 \} ,$$

where λ is a positive regularization parameter reflecting the relative value of information with respect to costs. Such a reward function reflects the notion that negative rewards should

result in taking no action. However, it can be shown that this reward function is not submodular and, as such, does not readily come with provable performance guarantees.

However, the related reward function, denoted as *penalized-MI* is

$$f(\mathcal{S}) = I(\mathbf{X}; Z^{\mathcal{S}}) - \lambda c(\mathcal{S}) ,$$

where $c(\mathcal{S})$ has been defined in (1). This set function is submodular, hence Eq. 4 holds with $\theta = \lambda c_{\max}$ and c_{\max} being the worst-case cost. However, this reward is non-monotone, *i.e.*, the increment function

$$\rho_{\mathcal{S}}(j) = f(\mathcal{S} \cup \{j\}) - f(\mathcal{S}) = I(\mathbf{X}; Z^j | Z^{\mathcal{S}}) - \lambda c(j) ,$$

can conceivably incur *negative* increments and the bounds for the sequential case (*cf.* [10]) no longer apply.

The optimal solution, however, can be shown (via contradiction) to contain only non-negative increments. Furthermore, via the introduction of k_t auxiliary measurements for each measurement set \mathcal{V}^t , the increments of the greedy solution will be non-negative as well.

2.2. Cost Proportional to Expected Information Reward

Typically, one would expend resources proportionate to the informational value of a measurement given the information that is already gathered. Here, we consider a proportional cost structure such that under certain conditions penalized MI is both submodular and monotone. Hence, the sequential guarantees of the greedy heuristic hold.

Consider costs which are proportional to the expected information gain of a measurement.

$$c_{\mathcal{S}}(j) \triangleq r_j I(\mathbf{X}; Z^j | Z^{\mathcal{S}}) ,$$

where r_j is a nonnegative constant associated to each measurement. The increment function thus becomes

$$\rho_{\mathcal{S}}(j) = I(\mathbf{X}; Z^j | Z^{\mathcal{S}}) - \lambda c_{\mathcal{S}}(j) = (1 - \lambda r_j) I(\mathbf{X}; Z^j | Z^{\mathcal{S}}) ,$$

which can be easily shown to be submodular. An additional requirement that ensures well-posedness is:

$$\begin{aligned} f(\mathcal{S} \cup \{j\}) &= \rho_{\mathcal{S}}(j) + \rho_{\emptyset}(\mathcal{S}) = \rho_j(\mathcal{S}) + \rho_{\emptyset}(j) \Leftrightarrow \\ c_{\mathcal{S}}(j) + c_{\emptyset}(\mathcal{S}) &= c_j(\mathcal{S}) + c_{\emptyset}(j) , \end{aligned}$$

where $c_{\emptyset}(j) = c(j)$ represents the cost of measurement Z_j . Consequently, we must have that $r_j = r, \forall j$. If, furthermore, $r \leq \frac{1}{\lambda}$ and $c(j) \leq r I(\mathbf{X}; Z^j | Z^{\mathcal{S}}), \forall j \in \mathcal{V}$ and $\mathcal{S} \subseteq \mathcal{V}$, the reward function is non-decreasing as well. Therefore, the batch and sequential guarantees in [4] and [10] will apply, respectively.

3. RESULTS

Here we present an experiment whose purpose is to demonstrate the following counter-intuitive result. If one decomposes the penalized reward into the information gain,

$I(\mathbf{X}; Z^{\mathcal{S}})$, and the cost $c(\mathcal{S})$ it is possible for greedy selection of penalized rewards to achieve a higher information gain than greedy selection of non-penalized rewards.

We consider the following linear state-space model:

$$\begin{aligned} X_{t+1} &= F X_t + U_t \\ Z_{t+1}^j &= H_{t+1}^j X_{t+1} + W_t , \end{aligned}$$

where $t = \{1, \dots, T\}, j = \{1, 2\}$, F captures linear dynamics, $U_t \sim \mathcal{N}(0, Q)$ is driving noise, and $W_t \sim \mathcal{N}(0, \sigma_t^2)$ is measurement noise. We set $\lambda = 0.5$ and consider $T = 10$ time points. For odd time points, we obtain two measurements which extract the position in x and y . For even time points, only the position in x is extracted. In addition, the variances of the noisy measurements are: $(\sigma_t^1)^2 = 16, (\sigma_t^2)^2 = 64$ for odd time points, while $(\sigma_t^1)^2 = 1$ for even time points. In addition, we set the costs such that $c(1) = 0.5, c(2) = 0.05, \forall t \in \{1, \dots, T\}$.

Let \mathcal{G}^I denote the set obtained using greedy MI as the reward function and \mathcal{G}^P denote the set obtained using greedy PMI. In Figure 1(top), we see that $I(\mathbf{X}; Z^{\mathcal{G}^I})$ is *lower* than $I(\mathbf{X}; Z^{\mathcal{G}^P})$, despite the fact that the \mathcal{G}^P selections incorporate penalties. The reason is twofold, first the greedy heuristic using MI as the reward selects measurements without regard to costs and second, the costs are structured in a way that the greedy choices for PMI prefer to measure the y -position of a latent variable when only the measurement of the x -position is available in neighboring nodes. This is not to say that there is a general method for adapting cost structures so that greedy PMI outperforms MI, rather that the two, in general, are not comparable and that it is quite possible (as our example shows) that PMI may yield better information rewards than pure MI.

Figure 1(middle) compares $I(\mathbf{X}; Z^{\mathcal{G}^I}) - \lambda c(\mathcal{G}^I)$ with $I(\mathbf{X}; Z^{\mathcal{G}^P}) - \lambda c(\mathcal{G}^P)$. Not surprisingly, the latter is superior to the former. We also see that the monotone behavior induced by the introduction of auxiliary variables results in non-negative increments \mathcal{G}^P . Lastly, Figure 1(bottom) shows that the posterior entropy of each latent variable is lower for \mathcal{G}^P as compared to \mathcal{G}^I .

4. DISCUSSION

We have shown that when costs are proportional to expected information gain and that the constant of proportionality satisfies certain conditions, penalized information rewards can be made to be monotone and submodular. The consequence of which is that the batch and sequential guarantees for greedy selection as compared to the optimal selection described in [4, 10] hold. Additionally, we presented simulations to illustrate the difference in unpenalized versus penalized measurement selection in an example tracking problem which illustrates the difficulty of directly comparing penalized to non-penalized rewards.

5. REFERENCES

- [1] E. Ertin, J. W. Fisher III, and L. C. Potter. Maximum Mutual Information Principle for Dynamic Sensor Query Problems. In Feng Zhao, editor, *Proceedings of the 2nd International Workshop on Information Processing in Sensor Networks*, pages 405–416. Springer-Verlag, Feb 2003.
- [2] S. Fujishige. *Submodular Functions and Optimization*. Annals of discrete mathematics. Elsevier, 2005.
- [3] A. Krause and C. Guestrin. A Note on the Budgeted Maximization of Submodular Functions. Technical Report CMU-CALD-05-103, Carnegie Mellon University, June 2005.
- [4] A. Krause and C. Guestrin. Near-optimal Nonmyopic Value of Information in Graphical Models. In *Conference on Uncertainty in Artificial Intelligence (UAI)*, July 2005.
- [5] A. Krause and C. Guestrin. Submodularity and its Applications in Optimized Information Gathering. *ACM Transactions on Intelligent Systems and Technology*, 2(4), 2011.
- [6] J. Lee, V. S. Mirrokni, V. Nagarajan, and M. Sviridenko. Maximizing Nonmonotone Submodular Functions under Matroid or Knapsack Constraints. *SIAM Journal of Discrete Mathematics*, 23(4):2053–2078, 2010.
- [7] G. L. Nemhauser, L. A. Wolsey, and M. L. Fisher. An Analysis of Approximations for Maximizing Submodular Set Functions - I. *Mathematical Programming*, 14(1):265–294, 1978.
- [8] J. L. Williams. *Information Theoretic Sensor Management*. PhD thesis, Massachusetts Institute of Technology, February 2007.
- [9] J. L. Williams, J. W. Fisher III, and A. S. Willsky. Approximate Dynamic Programming for Communication-Constrained Sensor Network Management. *IEEE Transactions on Signal Processing*, 55(8):3995–4003, August 2007.
- [10] J. L. Williams, J. W. Fisher III, and A. S. Willsky. Performance Guarantees for Information Theoretic Active Inference. In *Proc. Eleventh International Conference on Artificial Intelligence and Statistics*, March 2007.
- [11] F. Zhao, J. Shin, and J. Reich. Information-driven Dynamic Sensor Collaboration. *Signal Processing Magazine, IEEE*, 19(2):61–72, March 2002.




Fig. 1. Owing to the particular cost structure (top) \mathcal{G}^I results in a lower information reward than \mathcal{G}^P . (middle) \mathcal{G}^P outperforms \mathcal{G}^I (as expected) in terms of penalized increments. The introduction of auxiliary measurements yields non-negative increments. (bottom) The posterior entropy of each hidden variable is consistently lower for \mathcal{G}^P as compared to \mathcal{G}^I .