Theoretical Guarantees on Penalized Information Gathering

Georgios Papachristoudis and John W. Fisher III

Massachusetts Institute of Technology (MIT)

Abstract

We consider the problem of selecting measurements from an available subset for performing inference in graphical models. The measure of utility is an information reward discounted by resource costs. Optimal subset selection is known to be combinatorially complex and generally intractable even for sets of moderate size. However, it has been shown that greedy selection using unpenalized conditional mutual information rewards achieves performance within a factor of the optimal. Here we provide conditions under which cost-penalized mutual information may achieve similar guarantees.

Motivation

- Previous guarantees extend trivially to homogenous measurement costs.
- Consequently, we consider heterogenous measurement costs, specifically we are interested in cases where:
  1. costs are reflected as additive penalties to information rewards,
  2. differ across measurements, and
  3. information costs may vary, *e.g*, due to changing supply/demand for sensing/measurement assets.

Problem Statement

Notation

- $\mathbf{X} = \{X_1, \dots, X_n\}$ denotes n latent variables that are the focus of inference.
- $\mathbf{Z} = \{Z_1, \dots, Z_n\}$ denotes n measurement vectors.
- Each Z_t is comprised of N_t measurements corresponding to variable X_t.
- $\mathcal{V}^t = \{1, \dots, N_t\}$ indicate measurement indices, referred to as *observation sets*.
- $Z_i \perp\!\!\!\perp Z_j \mid \mathbf{X}$: we assume that measurements are statistically independent conditioned on $\mathbf{X}$.

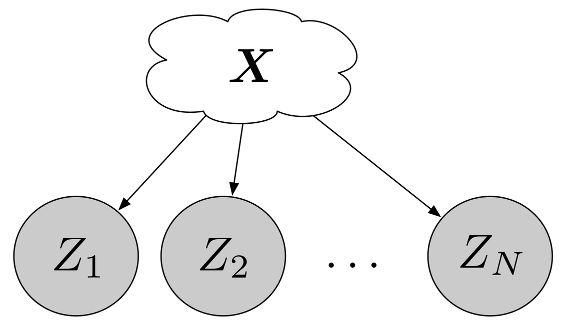

Batch Setting

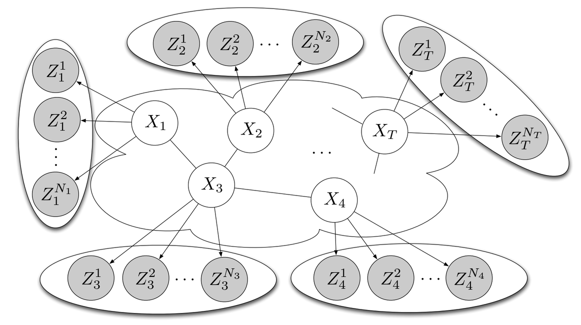

Sequential Setting

Batch setting

- $\mathbf{X}$ is treated as monolithic hidden variable.
- $\mathbf{Z}$ is treated as a uniform collection of measurements that are all simultaneously available.
- *Goal*: Find the best subset up to size K, where $K < N$ and $N = \sum_{t=1}^{n} N_t$: $\mathcal{O} = \arg\max_{|\mathcal{S}| \le K} f(\mathcal{S})$.

Sequential setting

- N_t measurements for each hidden variable X_t.
- *Visit walk*: define the M-length visit walk as the order $\{w_1, \dots, w_M\}$ in which we visit observation sets $\mathcal{V}^t$ during a selection process.
- *Goal*: Find the best subset up to k_1 measurements from set $\mathcal{V}^1$, …, k_n measurements from set $\mathcal{V}^n$: $\mathcal{O} = \arg\max_{|\mathcal{S}^1| \le k_1, \dots, |\mathcal{S}^n| \le k_n} f(\mathcal{S})$ where $\mathcal{S} = \bigcup_{t=1}^{n} \mathcal{S}^t$ and $\mathcal{S}^i \cap \mathcal{S}^j = \emptyset, \forall i \ne j$.

Reward function:

- $f : 2^{\mathcal{V}} \to \mathbb{R}$: a set function that captures the value of sensing actions.

Cost function:

- $c : 2^{\mathcal{V}} \to \mathbb{R}_+$: a nonnegative set function that quanitfies the cost of a subset and where all costs are assumed to be additive over the elements of the subset.

Submodularity

Given a finite set $\mathcal{V}$, a real-valued function f on the set of subsets of $\mathcal{V}$ is *submodular* [Fujishige, 05] if

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

- *Set increment function*: $\rho_{\mathcal{S}}(j) \triangleq f(\mathcal{S} \cup \{j\}) - f(\mathcal{S})$.
- A real-valued function is submodular if: $\rho_{\mathcal{A}}(j) \ge \rho_{\mathcal{B}}(j), \forall \mathcal{A} \subseteq \mathcal{B} \subseteq \mathcal{V}$ and $j \notin \mathcal{B}$.
- *Submodularity* captures the notion of "diminishing returns", that is, a measurement incurs higher rewards when little prior information is available.

Monotonicity

A real-valued f is *monotone* if $f(\mathcal{A}) \le f(\mathcal{B}), \forall \mathcal{A} \subseteq \mathcal{B}$ or $\rho_{\mathcal{S}}(j) \ge 0, \forall j \in \mathcal{V}, \mathcal{S} \subseteq \mathcal{V}$.

Greedy Selection

Greedy selection in the *batch* and *sequential* settings have different behavior owing to which measurements are available for selection at any given iteration.

Batch setting	Sequential setting
$g_j = \arg\max_{u \in \mathcal{V} \setminus \mathcal{G}^{j-1}} \rho_{\mathcal{G}^{j-1}}(u)$	$g_j = \arg\max_{u \in \mathcal{V}^{w_j} \setminus \mathcal{G}^{j-1}} \rho_{\mathcal{G}^{j-1}}(u)$

The *batch* setting chooses from among **all** measurements conditioned on previous selections while the *sequential* setting is restricted to only those available at the current node in the *visit walk*.

Mutual Information (MI) Rewards and Conditional Increments

Reward function: $f(\mathcal{S}) = I(\mathbf{X}; Z^{\mathcal{S}})$, Increment function: $\rho_{\mathcal{S}}(j) = I(\mathbf{X}; Z^j \mid Z^{\mathcal{S}})$

- monotone: $I(\mathbf{X}; Z^j \mid Z^{\mathcal{S}}) \ge 0, \forall j$ and $\mathcal{S} \subseteq \mathcal{V}$
- *submodular*: $\rho_{\mathcal{A}}(j) = I(\mathbf{X}; Z^j \mid Z^{\mathcal{A}}) \ge I(\mathbf{X}; Z^j \mid Z^{\mathcal{B}}) = \rho_{\mathcal{B}}(j), \forall \mathcal{A} \subseteq \mathcal{B}$.

Related Work

- [Nemhauser et al., 78I]: Greedy selection over monotone submodular functions for the batch setting is guaranteed to be at least $1 - 1/e \approx 0.63$ of the optimal reward.
- [Krause et al., 05]: Under the assumption of conditionally independent measurements, conditional mutual information is a monotone submodular function.
- [Nemhauser et al., 78II]: Greedy methods are within a bound of $\frac{1}{P+1}$ for monotone submodular functions under P matroid constraints where all measurements should be available at every iteration. The sequential setting is equivalent to a problem where each observation set forms a *uniform* matroid constraint $\mathcal{U}_{N_t, k_t}$, but only considers a limited class of visit paths (*i.e.*, those that observation sets cannot be revisited once we depart from them).
- [Lee et al., 10]: For non-negative non-monotone submodular functions greedy can be bounded by $\frac{1}{P+2+\frac{1}{P}+\epsilon}$ under P matroid constraints.
- [Williams et al., 07]: Proves a sharp bound of $1/2$ for greedy versus optimal selection for an arbitrary path in the sequential setting. Additional computable on-line bounds are available. *Remark*: The entire problem structure (latent or observation sets) need not be fully specified.

Nonmonotone Rewards

Let $\mathcal{G}$ and $\mathcal{O}$ denote the greedy and optimal subsets, respectively, and suppose that

- $\rho_{\mathcal{S}}(j) \ge -\theta, \forall \mathcal{S} \subseteq \mathcal{V}, j \in \mathcal{V} \setminus \mathcal{S}$, where $\theta \ge 0$ (*i.e.*, $-\theta$ is the worst case negative reward).
- Greedy selection terminates in $L \le K$ steps.

Batch setting

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

Sequential setting

$$f(\mathcal{O}) \le 2f(\mathcal{G}) + K\theta$$

Comment: For nonmonotone rewards, bounds become less useful as K and L grow for both batch and sequential settings.

Penalized Rewards and Modified Costs

Question: Can we impose additional constraints on either the reward or selection structure to improve the bounds?

Penalized Rewards: A natural choice is to penalize the information reward by the cost of acquiring a set of measurements via some positive conversion factor λ, *i.e.*,

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

While this is *submodular*, it is not *monotone* and we are left with the bounds above.

Thresholded Rewards: One could only select those measurements for which the reward is strictly positive, *i.e.*,

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

This is neither *monotone* nor *submodular* and as such we cannot use the bounds above to make comparisons.

Auxiliary Measurements: One could modify the *selection* structure to include k_t measurements for each set $\mathcal{V}^t$ for which the cost and information rewards are both zero.

- Behaves as the thresholded case, but is *submodular* and *nonmonotone*, *i.e.*, the bounds above apply.
- Interestingly, it can be shown that the greedy and optimal selection will have no negative incremental rewards and, as such, all rewards will be nonnegative as well.

Proportional Costs: An intuitive choice is to only pay costs which are proportional to the (informational) reward, captured by

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

where r_j is a nonnegative constant specific to each measurement.

- **well-posedness constraint:** The relation $\rho_{\mathcal{S}}(j) + \rho_{\emptyset}(\mathcal{S}) = \rho_j(\mathcal{S}) + \rho_{\emptyset}(j)$ must hold to preserve *well-posedness*, so $r_j = r, \forall j$.
- **nonmonotone constraint:** $c(j) \le rI(\mathbf{X}; Z^j \mid Z^{\mathcal{S}}), \forall j \in \mathcal{V}, \mathcal{S} \subseteq \mathcal{V}$ ensures *monotonicity*.

These conditions on proportional costs are sufficient such that the bounds of [Krause et al., 05] and [Williams et al., 07] apply and, as such, provide an extension to heterogeneous cost scenarios.

Simple Filtering Example

Consider the standard linear state-space model:

$$X_{t+1} = FX_t + U_t$$
$$Z_{t+1}^j = H_{t+1}^j X_{t+1} + W_t^j,$$

where $t \in \{1, \dots, T\}$, $j \in \{1, 2\}$, F desribes the state dynamics and $U_t \sim \mathcal{N}(0, Q)$, $W_t^j \sim \mathcal{N}(0, \sigma_{tj}^2)$ are driving and measurement noise, respectively.

Measurement Structure

- At odd time points, noisy measurements of the x- and y-positions, respectively, are available with $\sigma_{tx}^2 = 16$ and $\sigma_{ty}^2 = 64$.
- At even time points, only one noisy measurement of the x-position is available with $\sigma_{tx}^2 = 1$.
- The cost of the measurements for the x- and y-positions are $c = 0.5$ and $c = 0.05$, respectively.
- Let $\mathcal{G}^I, \mathcal{G}^P$ denote the sets obtained using MI and penalized MI, respectively, as the reward for selecting measurements.

Comparisons of Rewards

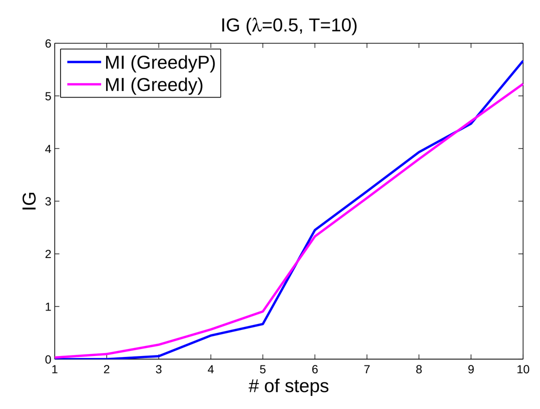

Information Reward

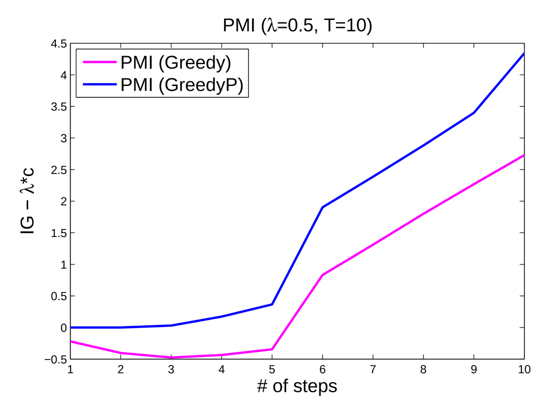

Penalized Reward

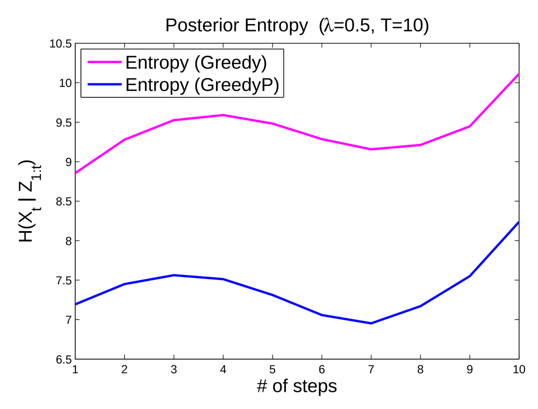

Posterior Uncertainty

Information Reward: Here 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 the measurement of the x-position is available in neighboring nodes. This is to say that depending on the cost structure the PMI reward may yield better information rewards than pure MI.

Penalized Reward: Comparison between $I(\mathbf{X}; Z^{\mathcal{G}^I}) - \lambda c(\mathcal{G}^I)$ and $I(\mathbf{X}; Z^{\mathcal{G}^P}) - \lambda c(\mathcal{G}^P)$. 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 in $\mathcal{G}^P$.

Posterior Entropy: The posterior entropy of each hidden variable is consistently lower for $\mathcal{G}^P$ as compared to $\mathcal{G}^I$.

References

- [Ertin et al., 03] E. Ertin, J. W. Fisher III, and L. C. Potter. Maximum Mutual Information Principle for Dynamic Sensor Query Problems. *Proceedings of the 2nd International Workshop on Information Processing in Sensor Networks*, pages 405-416. Springer-Verlag, Feb 2003.
- [Fujishige, 05] S. Fujishige. *Submodular Functions and Optimization*. Annals of discrete mathematics. Elsevier, 2005.
- [Krause et al., 05] Near-optimal Nonmyopic Value of Information in Graphical Models. *In Conference on Uncertainty in Artificial Intelligence (UAI)*, July 2005.
- [Krause et al., 11] Submodularity and its Applications in Optimized Information Gathering. *ACM Transactions on Intelligent Systems and Technology*, 2(4), 2011.
- [Lee et al., 10] 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.
- [Nemhauser et al., 78I] G. L. Nemhauser, L. A. Wolsey, and M. L. Fisher. An Analysis of Approximations for Maximizing Submodular Set Functions - I. *Mathematical Programming*, 14:265-294, 1978.
- [Nemhauser et al., 78II] G. L. Nemhauser, L. A. Wolsey, and M. L. Fisher. An Analysis of Approximations for Maximizing Submodular Set Functions - II. *Mathematical Programming*, 8:73-87, 1978.
- [Williams et al., 07] 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.

mail: geopapa@mit.edu, fisher@csail.mit.edu

www: http://groups.csail.mit.edu/vision/sli/