Online Submodular Set Cover, Ranking, and Repeated Active Learning

Andrew Guillory, Jeff Bilmes
University of Washington
guillory@cs.washington.edu, bilmes@ee.washington.edu

Contributions

- A new problem, *online submodular set cover*, with applications to ranking and repeated active learning
- A low regret algorithm
- Extensions to handle multiple objectives per round, partial information, and context
- Encouraging experimental results

Motivation

Online Ranking: At each round, learner produces an ordered list, suffers loss or receives reward.
Example: search result ranking

We are specifically interested in applications where the loss is the number of items needed to achieve some objective.
Example: The cost is the the number of pages the user needs to view to find the information they desire.

Repeated Active Learning is an interesting special case where the items are questions or tests. *Example: diagnosis.*

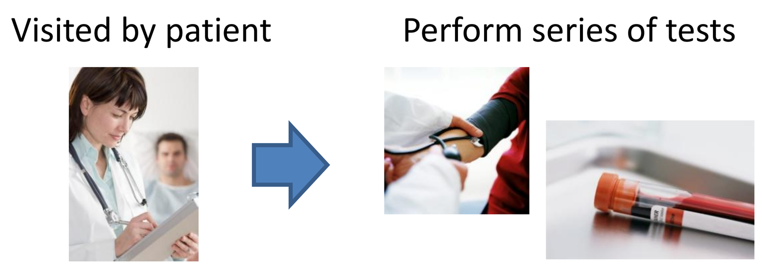

Here a reasonable loss is the number of the tests needed to make a confident diagnosis.

Background

Set function $F(S)$ defined over a ground set V is called *submodular* if for $A \subseteq B \subseteq V \setminus \{v\}$

$$F(A + v) - F(A) \geq F(B + v) - F(B)$$

Many natural objectives measuring information, influence, and coverage are submodular. $F(S)$ is *monotone* if for $A \subseteq B$, $F(A) \leq F(B)$ and *normalized* if $F(\emptyset) = 0$.

Submodular set cover is the problem of minimizing $|S|$ s.t. $F(S) \geq 1$ where F is submodular, monotone, normalized.

Online Submodular Set Cover

We propose *online submodular set cover* as a model for ranking and repeated active learning applications.

- At round t we pick a sequence $S^t = v_1, v_2, \ldots v_n$.
- A submodular, monotone, normalized F^t is revealed.
- We suffer loss $\ell(F^t, S^t)$ equal to the cover time of F^t:
  $$\ell(F^t, S^t) \triangleq \min\big(\{n\} \cup \{i : F^t(S_i^t) \geq 1\}_i\big)$$
  where $S_i^t \triangleq \bigcup_{j \leq i} \{v_j^t\}$
- Goal: minimize total loss $\sum_{i=1}^{T} \ell(F^t, S^t)$

Example: $F^t(S)$ proportional to the number of candidate diseases eliminated by performing tests S on patient t.

Related Work

Related to work by Streeter and Golovin (2008)

- Online submodular maximization: goal at round t is to select a set S^t with $|S^t| \leq k$ to maximize $F^t(S^t)$.
- Online min-sum submodular set cover: goal at round t is to select a sequence of items S^t to minimize
  $$\hat{\ell}(F^t, S^t) \triangleq \sum_{i=0}^{n} \max(1 - F^t(S_i^t), 0).$$
  Min-sum submodular set cover penalizes $1 - F^t(S_i^t)$ where submodular set cover uses $I(F^t(S_i^t) < 1)$.

Online submodular set cover makes more sense when the goal is to achieve a hard objective with minimal cost.

Offline Analysis

Our algorithm for the online problem is based on an algorithm of Azar and Gamzu (2011) for an offline problem, *ranking with submodular valuations*. There the goal is to compute the sequence S^* minimizing $\sum_{i=1}^{T} w_i \ell(F^t, S^*)$
Azar and Gamzu show the following greedy method approximates S^*. It uses a relative gain term

$$\delta(F^t, S, v) \triangleq \begin{cases} \min\left(\frac{F^t(S+v) - F^t(S)}{1 - F^t(S)}, 1\right) & \text{if } F(S) < 1 \\ 0 & \text{otherwise} \end{cases}$$

Offline Adaptive Residual Algorithm

Input: Objectives $F^1, F^2, \ldots F^T$
Output: Sequence $S_1 \subset S_2 \subset \ldots S_n$
$S_0 \leftarrow \emptyset$
for $i \leftarrow 1 \ldots n$ **do**
  $v \leftarrow \operatorname{argmax}_{v \in V} \sum_t \delta(F^t, S_{i-1}, v)$
  $S_i \leftarrow S_{i-1} + v$
end for

For our online problem, we needed a stronger version of the analysis of Azar and Gamzu. This analysis shows the algorithm works even with additive noise.

Theorem
Let $S = (v_1, v_2, \ldots v_n)$ be any sequence for which
$$\sum_t \delta(F^t, S_{i-1}, v_i) + R_i \geq \max_{v \in V} \sum_t \delta(F^t, S_{i-1}, v)$$
Then $\sum_t \ell(F^t, S^t) \leq 4(\ln 1/\epsilon + 2) \sum_t \ell(F^t, S^*) + n \sum_i R_i$

Online Analysis

Our online algorithm lifts the offline algorithm to the online setting, similar to Streeter and Golovin (2008).

Online Adaptive Residual Algorithm

Input: Integer T
Initialize n online learning algorithms $E_1, \ldots E_n$
for $t = 1 \to T$ **do**
  $\forall i \in 1 \ldots n$ predict v_i^t with E_i
  $S^t \leftarrow (v_1^t, \ldots v_n^t)$
  Receive F^t, pay loss $l(F^t, S^t)$
  For E_i, $\ell^t(v) \leftarrow (1 - \delta(F^t, S_{i-1}^t, v))$
end for

Theorem
Assume E_i has regret $\mathbb{E}[R_i] \leq \sqrt{T \ln n}$. The Online Adaptive Residual Algorithm has α-regret $\mathbb{E}[R_\alpha] \leq n^2\sqrt{T \ln n}$ for $\alpha = 4(\ln(1/\epsilon) + 2)$

The proof makes use of our strengthened analysis of the offline algorithm and the particular way in which the loss for the black box prediction algorithms are constructed.

Extensions

Our proposed online algorithm is easy to modify and analyze for many variations of the problem setting

- **Truncated Loss:** a version with a truncated loss with parameter k
  $$\ell^k(F^t, S^t) \triangleq \min\big(\{k\} \cup \{|S_i^t| : F^t(S_i^t) \geq 1\}\big)$$
- **Multiple Objectives per Round:** a variation where at each round we receive a batch of objectives $F_1^t, F_2^t, \ldots F_m^t$ and incur loss $\sum_{i=1}^{m} \ell(F_i^t, S^t)$.
- **Partial Information Setting:** a version where we only observe the sequence of objective function values $F^t(S_1^t), F^t(S_2^t), \ldots F^t(S_n^t)$
- **Partial Information Setting with Expert Advice:** a version where we also have access at time step t to item selections from a set of m experts.

Experiments

We compare our adaptive residual method to the cumulative greedy method of Streeter and Golovin (2008) which has approximation guarantees for online submodular maximization but *not* for online submodular set cover.

Synthetic Example

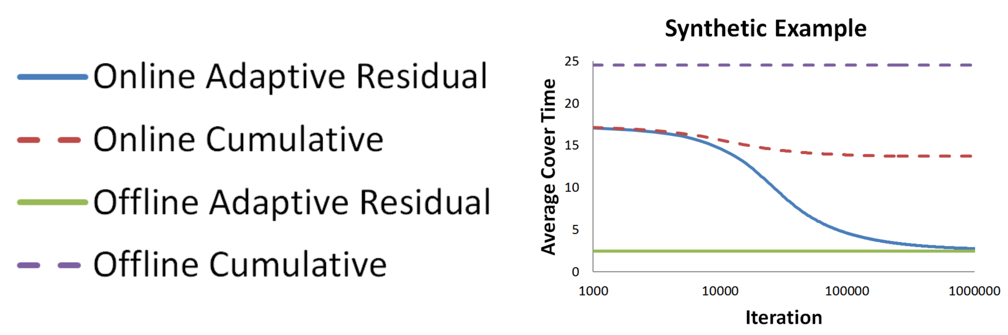

A synthetic example from Azar and Gamzu carries over to our online problem. Here $n = 25$. This shows that, in general, the average cover time of the cumulative greedy method is much worse than the adaptive method.

Movie Recommendation

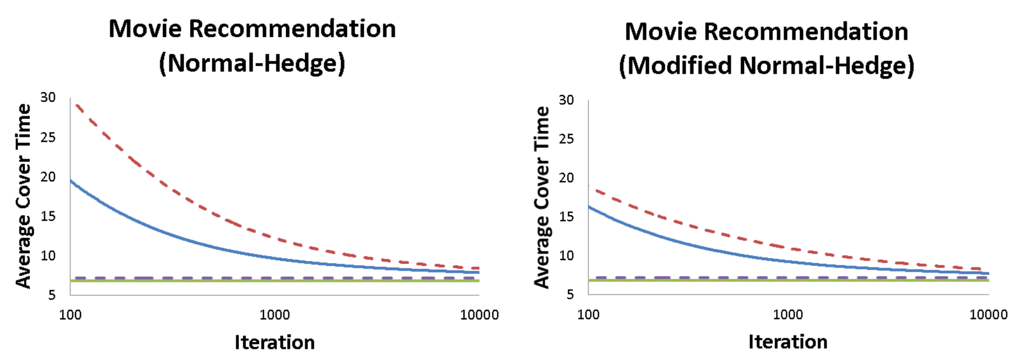

Here we consider choosing sequences of questions in order to quickly eliminate candidate movie recommendations The ground set V is a set of questions ("Do you want to watch something from the Drama genre?") and the set of movies is 11634 movies from Netflix's Watch Instantly Service. The adpative residual method again outperforms the cumulative greedy method. The difference is more dramatic when convergence is slower.