


Online Submodular Set Cover, Ranking, and Repeated Active Learning

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

Get search query



Display ranked results

[Cheap Auto Insurance Quotes | Get the Best Rates Fast](#)
www.autoinsurancequoteseasy.com/ [+]

Cheap **auto insurance** quotes even for young drivers. Find all the major vehicle insurance companies in one place and get free quotes instantly.

[GEICO | GEICO Car Insurance - Get online auto insurance quotes...](#)
www.geico.com/ [+]

Offers **auto** and motorcycle **insurance**. Request quotes online or by phone. File claims online.

[Manage Your Policy - Contact Us - Get A Quote - Make a Payment](#)

[State Farm - Car Insurance Quotes - Save on Auto Insurance](#)
www.statefarm.com/ [+]

State Farm offers more discounts to more drivers than any other **insurance** company. Get a free quote today and see how you can save on **auto**, home, life and ...

In this paper, loss is:

the number of items needed to achieve some coverage objective

Example: The cost at each round is the number of pages the user needs to view to deduce the complete information they desire.

Online Submodular Set Cover, Ranking, and Repeated Active Learning

Repeated Active Learning is an interesting special case where the list consists of questions to ask or tests to perform. *Example: diagnosis*

Visited by patient



Perform series of tests



Here a reasonable loss is the number of the tests we need to perform before we can make a accurate diagnosis.

Online Submodular Set Cover, Ranking, and Repeated Active Learning

For these applications we propose a new online learning problem we call **online submodular set cover**.

- At round t we pick a sequence $S^t = v_1, v_2, \dots, v_n$.
- A monotone, submodular objective F^t is then revealed.
- We pay cost equal to the cover time of F^t : the minimum value $c \in \{1, 2, \dots, n\}$ such that $F^t(\bigcup_{i=1}^c \{v_i\}) \geq 1$.

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

Related but not equivalent to online submodular maximization and online min-sum submodular set cover (Streeter and Golovin, 2008)

Online Submodular Set Cover, Ranking, and Repeated Active Learning

Our results:

- A low-regret algorithm for online submodular set cover, building on a recent offline algorithm of Azar and Gamzu.
- Extensions to handle multiple objectives, partial information, context.
- Experimental results on synthetic data and a movie recommendation repeated active learning problem.

- Online Adaptive Residual
- - Online Cumulative
- Offline Adaptive Residual
- - Offline Cumulative

**Movie Recommendation
(Modified Normal-Hedge)**

