Cornell University Library

We gratefully acknowledge support from the Simons Foundation and member institutions

arXiv.org > cs > arXiv:1305.2532

Search or Article-id

(Help | Advanced search)

All papers ▼   Go!

Computer Science > Learning

Learning Policies for Contextual Submodular Prediction

Stephane Ross, Jiaji Zhou, Yisong Yue, Debadeepta Dey, J. Andrew Bagnell

(Submitted on 11 May 2013)

Many prediction domains, such as ad placement, recommendation, trajectory prediction, and document summarization, require predicting a set or list of options. Such lists are often evaluated using submodular reward functions that measure both quality and diversity. We propose a simple, efficient, and provably near-optimal approach to optimizing such prediction problems based on no-regret learning. Our method leverages a surprising result from online submodular optimization: a single no-regret online learner can compete with an optimal sequence of predictions. Compared to previous work, which either learn a sequence of classifiers or rely on stronger assumptions such as realizability, we ensure both data-efficiency as well as performance guarantees in the fully agnostic setting. Experiments validate the efficiency and applicability of the approach on a wide range of problems including manipulator trajectory optimization, news recommendation and document summarization.

Submission history

From: Stephane Ross [view email]

[v1] Sat, 11 May 2013 18:09:52 GMT (169kb,D)

Which authors of this paper are endorsers?

Download:

- PDF
- Other formats

Current browse context:
cs.LG
< prev | next >
new | recent | 1305

Change to browse by:
cs
stat
    stat.ML

References & Citations

- NASA ADS

DBLP - CS Bibliography
listing | bibtex
Stéphane Ross
Jiaji Zhou
Yisong Yue
Debadeepta Dey
J. Andrew Bagnell

Bookmark (what is this?)