

Seminar Series

Mohsen Bayati

Associate Professor, Stanford University

Wednesday, September 28, 2016

12:30 pm

Tech M228



Online Decision-Making with High-Dimensional Covariates

ABSTRACT:

Growing availability of data has enabled decision-makers to tailor choices at the individual-level. This involves learning a model of decision rewards conditional on individual-specific covariates or features. Recently, "contextual bandits" have been introduced as a framework to study these online decision making problems. However, when the space of features is high-dimensional, existing literature only considers situations where features are generated in an adversarial fashion that leads to highly conservative performance guarantees -- regret bounds that scale by square-root of number of samples.

Motivated by medical decision making problems where stochastic features are more realistic, we introduce a new algorithm that relies on two sequentially updated LASSO estimators. One estimator (with a low-bias) is used to select a candidate subset of the decisions, next a more biased (but potentially more accurate) estimator is used to select the optimal decision. We prove that our algorithm achieves a regret that scales poly-logarithmically in the number of samples and features. The key step in our analysis is proving a new oracle inequality that guarantees the convergence of the LASSO estimator despite the non-i.i.d. data induced by the bandit policy.

We illustrate the practical relevance of the proposed algorithm by evaluating it on a warfarin dosing problem. A patient's optimal warfarin dosage depends on the patient's genetic profile and medical records; incorrect initial dosage may result in adverse consequences such as stroke or bleeding. We show that our algorithm outperforms existing bandit methods as well as physicians to correctly dose a majority of patients.

This is joint work with Hamsa Bastani.