Robust inference and local algorithms

Yishay Mansour, Tel Aviv University

Abstract

Robust inference is an extension of probabilistic inference, where some of the observations may be adversarially corrupted. We limit the adversarial corruption to a finite set of modification rules. We model robust inference as a zero-sum game between an adversary, who selects a modification rule, and a predictor, who wants to accurately predict the state of nature.

There are two variants of the model, one where the adversary needs to pick the modification rule in advance and one where the adversary can select the modification rule after observing the realized uncorrupted input. For both settings we derive efficient near optimal policy runs in polynomial time. Our efficient algorithms are based on methodologies for developing local computation algorithms.

We also consider a learning setting where the predictor receives a set of uncorrupted inputs and their classification. The predictor needs to select a hypothesis, from a known set of hypotheses, and is tested on inputs which the adversary corrupts. We show how to utilize an ERM oracle to derive a near optimal predictor strategy, namely, picking a hypothesis that minimizes the error on the corrupted test inputs.

Based on joint works with Uriel Feige, Aviad Rubinstein, Robert Schapira, Moshe Tennenholtz, and Shai Vardi.