Learning with low mutual information

Amir Yehudayoff and Krishna Manoorkar

Learning refers to the process of producing a hypothesis which best describes the available data (e.g. samples from some distribution). It can be defined in different ways depending on the specific setting. We use the formalism of PAC (probably approximately correct) learning . A learning algorithm A in this framework is a map that takes as input a list S of sample points in $\mathcal{X} \times \{0, 1\}$ and outputs a function $A(S) : \mathcal{X} \to \{0, 1\}$.

There are many interesting ways to quantify the quality of a learner, such as its computational complexity or its generalization capabilities. We focus on the amount of information the output $A(S)$ of A reveals on its inputs S. The standard way to measure information is via the mutual information $I(A; A(S))$.

The project's goal is to develop learning algorithms with low information complexity for specific scenarios. To achieve this goal, you will need to understand several deep definitions, like generalization, entropy, and the Littlestone dimension.

Prerequisites. Read [1] and follow the references within so that you understand the main definitions and ideas.

References

[1] Raef Bassily, Shay Moran, Ido Nachum, Jonathan Shafer, Amir Yehudayof (2018). *Learners that Use Little Information. Proceedings of Algorithmic Learning Theory, in PMLR* 83:25-55.