Beyond regression and classification: Structured Prediction with theoretical guarantees

Alessandro Rudi

January 2019

Standard ML techniques learn real valued functions. However nowadays many machine learning problems require to learn functions with structured output e.g. automatic text translation (string \rightarrow string), image captioning (image \rightarrow string), speech recognition (sound \rightarrow text), prediction on graphs, learning on manifolds, learning to rank documents, protein folding, etc. In this context current research [1] focuses on developing *ad-hoc* algorithms for each structured problem, often without theoretical guarantees.

Recently, [2] proposed a unifying theoretical framework to address in a systematic way structured prediction problems by introducing a novel learning strategy with strong statistical guarantees. In particular the derived methods are expressed in terms of an optimization problem in the output space. The approach is explained and further specified for many useful discrete settings in [3].

References

[1] Gükhan H. Bakir, Thomas Hofmann, Bernhard Schölkopf, Alexander J. Smola, Ben Taskar, and S. V. N. Vishwanathan. 2007. *Predicting Structured Data*. (Neural Information Processing). The MIT Press.

[2] Ciliberto, C., Rudi, A. and Rosasco, L. A Consistent Regularization Approach for Structured Prediction. In Advances in Neural Information Processing Systems 2016. (pp. 4412-4420). Barcelona

[3] Nowak-Vila, A., Bach, F., and Rudi, A. *Sharp analysis of learning with discrete losses.* arXiv preprint arXiv:1810.06839 (2018).