

EXACT GENERALIZATION GUARANTEES FOR (REGULARIZED) WASSERSTEIN DISTRIBUTIONALLY ROBUST OPTIMIZATION

Waïss Azizian

LJK, Université Grenoble Alpes , France
waiss.azizian@univ-grenoble-alpes.fr

Wasserstein distributionally robust estimators have emerged as powerful models for modeling decision-making under uncertainty. These estimators provide a new kind of generalization guarantee: the robust objective w.r.t. the training distribution is, with high probability, an exact upper bound on the true risk. However, existing guarantees suffer from the curse of dimensionality in their dependence on the number of samples, are restricted to specific setting, or require a vanishing uncertainty radius and lead to spurious error terms. We show that the generalisation bounds still hold on general classes of models with correct dependency on the dimension, and furthermore allow for a non-vanishing radius to cover distribution shifts at testing. We also prove that these guarantees carry over to the newly introduced regularized versions of Wasserstein distributionally robust problems.

Joint work with Franck Iutzeler (LJK, Université Grenoble Alpes) and Jérôme Malick (LJK, CNRS, Université Grenoble Alpes).