PhD proposal
Adaptive methods for safe machine learning

Université Nice Côte d'Azur
Lab. Jean Alexandre Dieudonné

Supervisors: Yassine Laguel, Samuel Vaiter
Duration : 3 years

Motivations In machine learning, safety concerns can be driven by various factors, stemming from data variability, algorithms instability, fairness and privacy constraints, or unpredictable changes in the environment. Distributionally Robust Optimization (DRO) [1] addresses these challenges by optimizing models over various potential data distributions, thus enhancing their robustness and reliability. Specifically, DRO expands the traditional Empirical Risk Minimization (ERM) paradigm by allowing train and test distribution to depart from each other. One seeks then to solve problems of the form

$$\min_{w} \max_{Q \in \mathcal{A}} \mathbb{E}_{(X,Y)\sim Q}[f(w, X, Y)]$$

where f denotes the loss of the learning problem, and the set $\mathcal{A}$, often called an ambiguity set, denotes an arbitrary neighborhood of the training distribution $\hat{P}_n = ((X_1, y_1), \ldots, (X_n, y_n))$. A key consideration in DRO is the trade-off between allowing for ambiguity in data distributions – *i.e.,* having $\mathcal{A}$ large – and the computational overhead this introduces [4].

Goals The aim of this PhD. is to develop simple automatic selection procedures for the ambiguity set $\mathcal{A}$, to effectively address safety constraints in data-driven learning contexts.

1. First, we aim to refine the coarse risk quantification in DRO methods. This issue originates from the use of overly broad distribution shift ranges, leading to uncertainty among ML practitioners about selecting the right distribution shift family. The goal is to develop automated methods for the selection of ambiguity sets $\mathcal{A}$ that adjust to the possible data shifts encountered, achieving a balance between model precision and the solvability these high-dimensional problems [2].
2. Second, we aim to develop stochastic first-order methods tailored to the structure of these min-max problems, particularly influenced by the geometry of the ambiguity set A. This approach may result in better efficiency compared to general min-max problems, as previously observed in [3]. Success in this area will lead to the release of open-source software and numerical benchmarks.

Candidate Profile This PhD proprosal is designed for a candidates with a strong background in applied Mathematics or computer science. Candidates with previous experience in optimization are appreciated, as is an interest in numerical simulations.

Application Procedure Send your CV and a copy of your last year transcripts to both Yassine Laguel (`yassine.laguel@univ-cotedazur.fr`) and Samuel Vaiter (`samuel.vaiter@cnrs.fr`)

References

[1] Lin, F. and Fang, X. and Gao, Z. (2022) Distributionally Robust Optimization: A review on theory and applications. Numerical Algebra, Control and Optimization.

[2] Laguel Y. and Pillutla, K. and Malick, J. and Harchaoui, Z. (2022) Federated learning with superquantile aggregation for heterogeneous data. Machine Learning Journal.

[3] Laguel Y. and Malick, J. and Harchaoui, Z. (2022) Superquantile-based learning: a direct approach using gradient-based optimization Journal of Signal Processing Systems.

[4] Levy, D. and Carmon, Y and Duchi, J.C. and Sidford, A. (2020) Large-Scale Methods for Distributionally Robust Optimization Neurips.