Learning Credal Ensembles via Distributionally Robust Optimization

Kaizheng Wang[1,2], Ghifari Adam Faza[3,4], Fabio Cuzzolin[5], Siu Lun Chau[2], David Moens[3,4], Hans Hallez[1]

[1]Department of Computer Science, KU Leuven
[2]College of Computing and Data Science, Nanyang Technological University
[3]Department of Mechanical Engineering, KU Leuven
[4]Flanders Make@KU Leuven
[5]School of Engineering, Computing, and Mathematics, Oxford Brookes University

Background and Motivation

Informative **uncertainty quantification** (UQ) in neural networks requires distinguishing:

- Aleatoric uncertainty (AU): inherent randomness in the data-generation process
- Epistemic uncertainty (EU): neural networks' limited knowledge about the true input-output relationship

Reliable estimation of EU is beneficial for tasks where utilizing model's predictive confidence is crucial:

- selective prediction
- learning to reject and defer
- **out-of-distribution (OOD) detection**
- Bayesian optimization

Background and Motivation

Credal predictors provide a principled framework for **quantifying predictive EU**:

- produce a credal set, i.e., a convex set of probabilistic predictions
- have been shown to improve model robustness across a range of settings
- However, SOTA methods primarily define EU as disagreement induced by random training initializations, which mainly reflects sensitivity to optimization randomness rather than uncertainty from more substantive sources

Proposed solution:

- *Formulate EU as the* ***disagreement among models that are trained under varying degrees of relaxation of the i.i.d. assumption*** *(i.e., allowing training and test distributions to differ)*

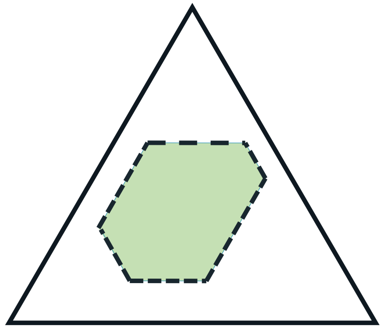

A credal prediction on 2D probability simplex

Learning Credal Ensembles via Distributionally Robust Optimization

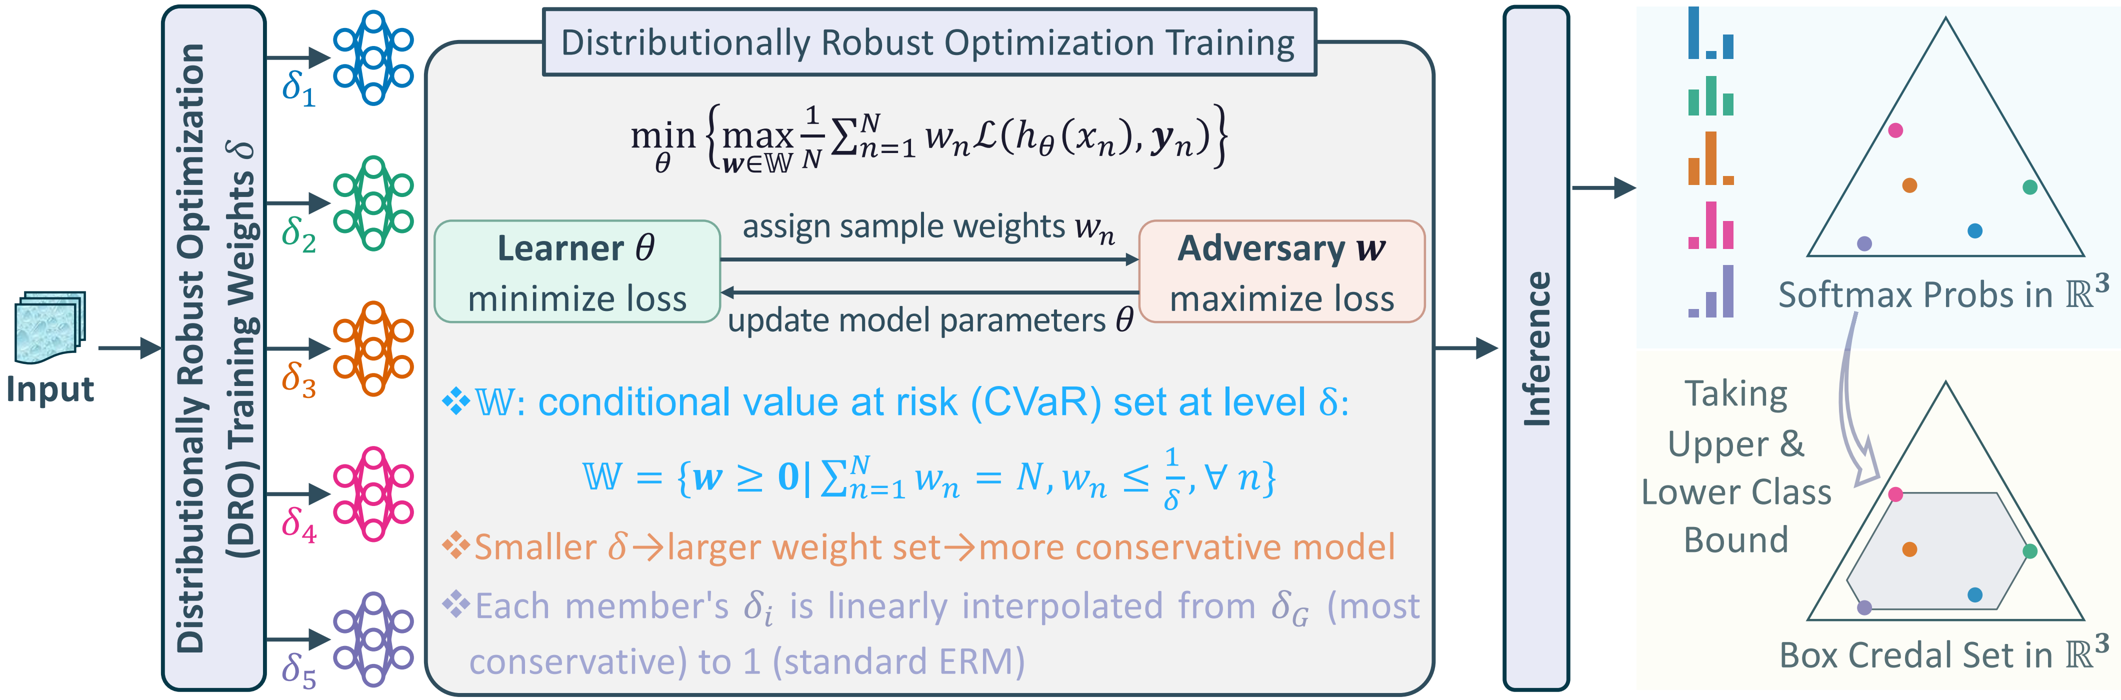

Experimental Validations

Empirically, our method consistently outperforms SOTA credal approaches on downstream tasks, including:

- out-of-distribution detection on several benchmarks
- selective classification in realistic medical settings

Practical & robustness:

- no architecture changes needed
- robust to hyperparameter choice δ_G & symmetric label noise

Thank You for Your Attention!