Quantifying Aleatoric and Epistemic Uncertainty
A Credal Approach

Paul Hofman, Yusuf Sale, and Eyke Hüllermeier

LMU Munich, MCML

Setting

Supervised classification with instance space $\mathcal{X}$ and label space $\mathcal{Y}$.
Training data
$\mathcal{D} = \{(x_i, y_i)\}_{i=1}^N \in (\mathcal{X} \times \mathcal{Y})^N$.
Hypothesis space
$\mathcal{H} = \{h : \mathcal{X} \to \mathbb{P}(\mathcal{Y})\}$.
Given an instance $\boldsymbol{x}$, we denote $h(\boldsymbol{x}) = \boldsymbol{\theta}$.

Types of Uncertainty [1]

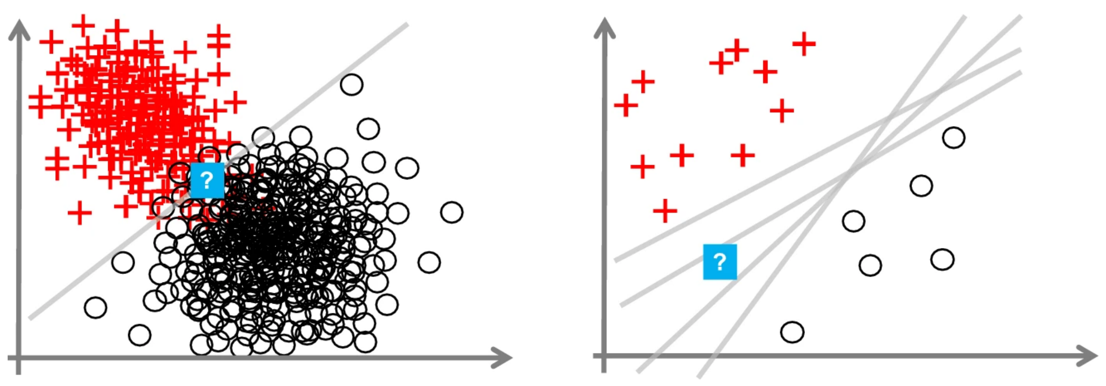

Aleatoric Uncertainty

Refers to the variability of the outcome due to the inherent randomness of the data and is therefore irreducible.

Epistemic Uncertainty

Refers to uncertainty caused by a lack of knowledge (epistemic state of the learner) and is reducible by adding information.

Uncertainty Representation

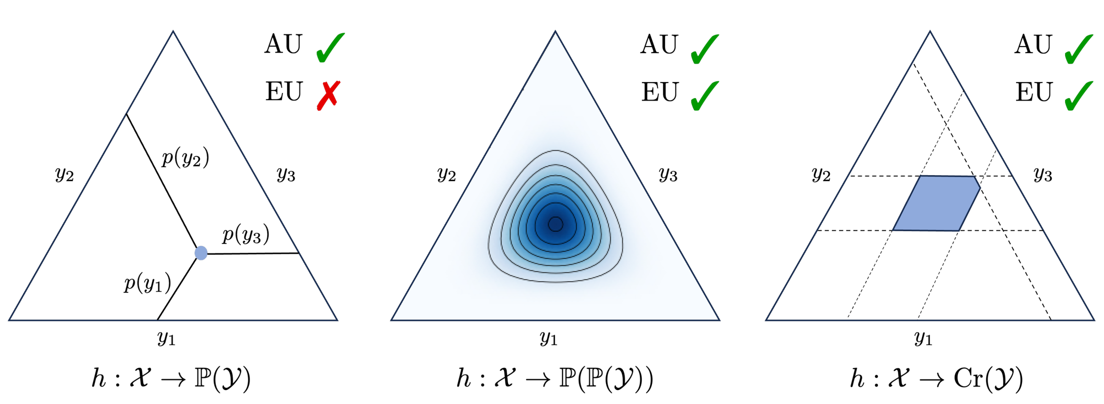

A credal set C is a (closed and convex) set of probability distributions on $\mathcal{Y}$.

$$\mathrm{TU}(C) = \mathrm{AU}(C) + \mathrm{EU}(C).$$

Various decompositions exist, i.a. based on

$$\mathrm{GH}(C) := \sum_{A \subseteq \mathcal{Y}} \mathrm{m}_Q(A) \log(|A|)$$

gen. Hartley [2] and Shannon entropy S.

$$\underbrace{S^*(C)}_{\mathrm{TU}} = \underbrace{(S^*(C) - \mathrm{GH}(C))}_{\mathrm{AU}} + \underbrace{\mathrm{GH}(C)}_{\mathrm{EU}},$$

$$\underbrace{S^*(C)}_{\mathrm{TU}} = \underbrace{S_*(C)}_{\mathrm{AU}} + \underbrace{(S^*(C) - S_*(C))}_{\mathrm{EU}},$$

where

$$S^*(C) := \max_{\boldsymbol{\theta} \in C} S(\boldsymbol{\theta}), \quad S_*(C) := \min_{\boldsymbol{\theta} \in C} S(\boldsymbol{\theta}).$$

A credal approach

Learning Credal Predictors

Relative likelihood given training data $\mathcal{D}$

$$\mathcal{L}_{\mathcal{H}}(h) := \frac{L(h)}{L(\hat{h})},$$

where $L(h) = \prod_{i=1}^N p(y_i \,|\, h, x)$ and $\hat{h}$ the (empirical) maximum likelihood predictor.

Given $\boldsymbol{x} \in \mathcal{X}$ and $Q \subseteq \mathcal{H}$ (plausible hypotheses, i.e. ensemble members):

$$C_\alpha := \{h(\boldsymbol{x}) \in Q \,|\, \mathcal{L}_{\mathcal{H}}(h) \geq \alpha\},$$

with $\alpha \in [0, 1]$.

Uncertainty Quantification

Epistemic uncertainty

$$\mathrm{EU}(C) := \max_{\boldsymbol{\theta}, \boldsymbol{\theta}' \in C} D_\ell(\boldsymbol{\theta}, \boldsymbol{\theta}'),$$

with the ℓ-divergence

$$D_\ell(\boldsymbol{\theta}, \boldsymbol{\theta}') := \mathbb{E}_{Y \sim \boldsymbol{\theta}} \{\ell(\boldsymbol{\theta}', Y) - \ell(\boldsymbol{\theta}, Y)\},$$

the excess loss of predicting $\boldsymbol{\theta}'$, while the ground truth is $\boldsymbol{\theta}$.

Aleatoric uncertainty

$$\underline{\mathrm{AU}}(C) := \min_{\boldsymbol{\theta} \in C} H_\ell(\boldsymbol{\theta}),$$

$$\overline{\mathrm{AU}}(C) := \max_{\boldsymbol{\theta} \in C} H_\ell(\boldsymbol{\theta}),$$

with H_ℓ the ℓ-entropy of $\boldsymbol{\theta}$ given by

$$H_\ell(\boldsymbol{\theta}) := \mathbb{E}_{Y \sim \boldsymbol{\theta}}\, \ell(\boldsymbol{\theta}, Y),$$

the irreducible loss of ground truth $\boldsymbol{\theta}$.

Here, ℓ can be any loss. We consider (strictly) proper scoring rules [3].

Loss	Aleatoric (upper\lower)	Epistemic
log	$\sup_{\boldsymbol{\theta} \in C} \backslash \inf_{\boldsymbol{\theta} \in C} S(\boldsymbol{\theta})$	$\max_{\boldsymbol{\theta}, \boldsymbol{\theta}' \in C} \mathrm{D}_{\mathrm{KL}}(\boldsymbol{\theta}' \,\|\, \boldsymbol{\theta})$
Brier	$\sup_{\boldsymbol{\theta} \in C} \backslash \inf_{\boldsymbol{\theta} \in C} 1 - \sum_{k=1}^K \theta_k^2$	$\max_{\boldsymbol{\theta}, \boldsymbol{\theta}' \in C} \sum_{k=1}^K (\theta_k - \theta'_k)^2$
spherical	$\sup_{\boldsymbol{\theta} \in C} \backslash \inf_{\boldsymbol{\theta} \in C} 1 - \|\|\boldsymbol{\theta}\|\|_2$	$\max_{\boldsymbol{\theta}, \boldsymbol{\theta}' \in C} \|\|\boldsymbol{\theta}'\|\|_2 - \sum_{k=1}^K \theta_k \theta'_k / \|\|\boldsymbol{\theta}'\|\|_2$
zero-one	$\sup_{\boldsymbol{\theta} \in C} \backslash \inf_{\boldsymbol{\theta} \in C} 1 - \max \theta_k$	$\max_{\boldsymbol{\theta}, \boldsymbol{\theta}' \in C} \max \theta'_k - \theta'_{k = \mathrm{argmax}\, \theta_k}$

Different losses account for different uncertainties [4].

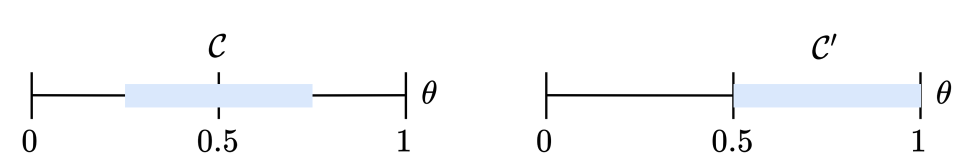

Left: decision uncertainty. Right: no decision uncertainty.

Results

Accuracy-Rejection Curves

Ensemble of 5 pre-trained ResNets fine-tuned on Food101. Upper aleatoric uncertainty is used as rejection criterion.

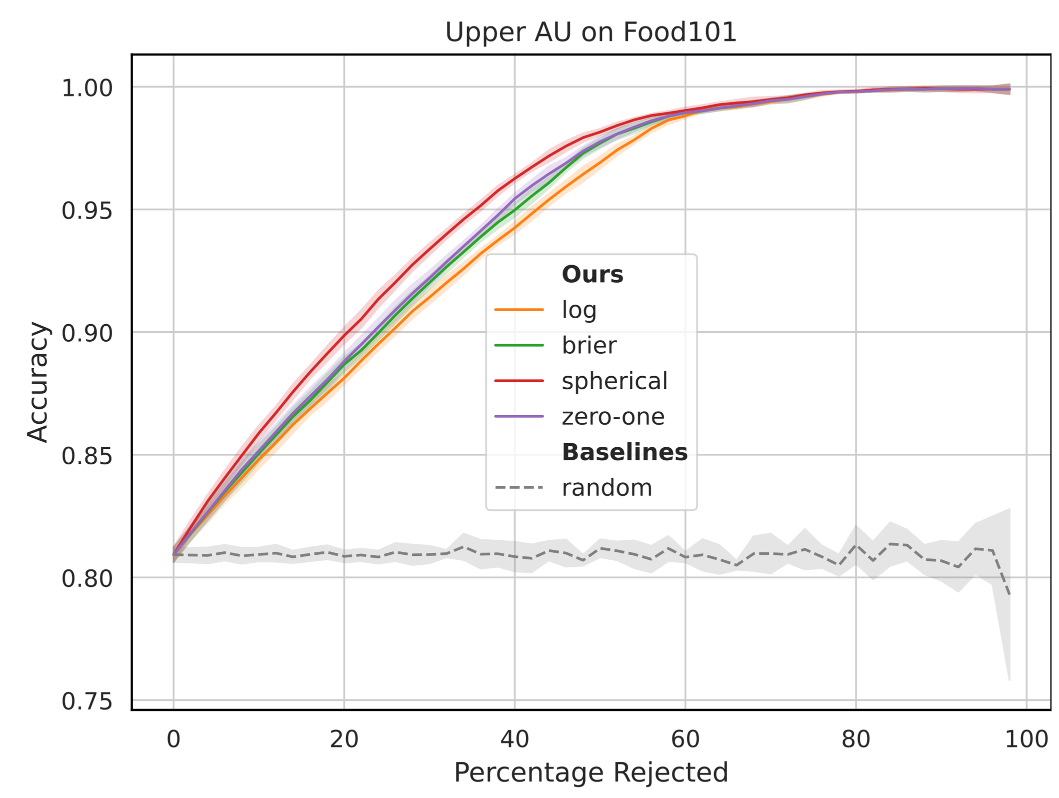

Out-of-Distribution Detection

Ensemble of 5 LeNets for FMNIST and 5 ResNets for Food101. OoD using epistemic uncertainty.

iD	OoD	log	Brier	spherical	zero-one	Hartley	entropy
FMNIST	MNIST	**0.871**± 0.017	0.812± 0.019	0.818± 0.019	0.742± 0.015	0.815± 0.019	0.826± 0.019
	KMNIST	**0.973**± 0.002	0.94± 0.004	0.946± 0.003	0.863± 0.006	0.944± 0.004	0.942± 0.003
Food101	SVHN	**0.700**± 0.04	0.572± 0.027	0.588± 0.038	0.669± 0.007	0.479± 0.023	0.681± 0.07
	CIFAR-100	**0.805**± 0.015	0.66± 0.016	0.681± 0.016	0.697± 0.008	0.576± 0.022	0.775± 0.021

Active Learning

Ensemble of 10 small MLPs. Initial training pool is 50 instances. Every round 50 new instances are acquired based on their epistemic uncertainty.

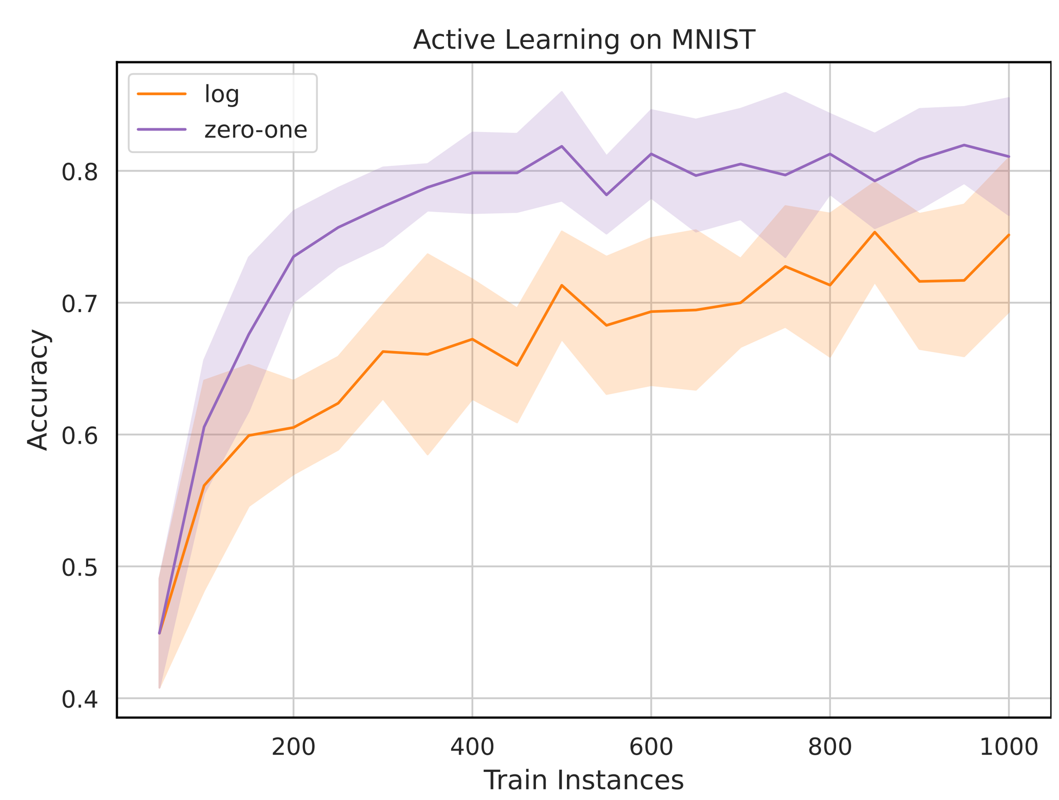

Future Work

- Can we offer guarantees on the credal set [5]?
- How does relaxing the convexity assumption affect the theoretical and empirical results?

[1] Eyke Hüllermeier and Willem Waegeman. "Aleatoric and epistemic uncertainty in machine learning: An introduction to concepts and methods". In: Machine Learning 110.3 (2021).
[2] Joaquín Abellán, George J. Klir, and Serafín Moral. "Disaggregated total uncertainty measure for credal sets". In: *International Journal of General Systems* 35.1 (2006).
[3] Tilmann Gneiting and Adrian Raftery. "Strictly Proper Scoring Rules, Prediction, and Estimation". In: *Journal of the American Statistical Association* 102.477 (2005).
[4] Paul Hofman, Yusuf Sale, and Eyke Hüllermeier. "Quantifying Aleatoric and Epistemic Uncertainty with Proper Scoring Rules". In: CoRR (2024).
[5] Alireza Javanmardi, David Stutz, and Eyke Hüllermeier. "Conformalized Credal Set Predictors". In: CoRR (2024).