Quantifying Classification Uncertainty using Regularized Evidential Neural Networks

Xujiang Zhao[1], Yuzhe Ou[1], Lance Kaplan[2], Feng Chen[1], Jin-Hee Cho[3]

[1]The University of Texas at Dallas, [2]US Army Research Lab, [3]Virginia Tech

Archive version available at: https://arxiv.org/pdf/1910.06864

Why is predicting uncertainty important?

Is it important to know:

- ✓ why we don’t know?
- ✓ how much we don’t know?

So how can we predict the uncertainty based on its root cause?

Would it really help for our decision making?

What types of uncertainty to model?

1

Epistemic uncertainty (a.k.a. model/parameter uncertainty)

- Measures what model doesn't know
- Due to **limited data and knowledge**

Aleatoric uncertainty (a.k.a. data uncertainty)

- Measures what you can't understand from the data
- Due to **randomness**

→ Probabilistic Uncertainty

2

***Vacuity* uncertainty** (a.k.a. ignorance)

- Measures uncertainty due to **a lack of evidence**

***Dissonance* uncertainty**

- Measures uncertainty due to **conflicting evidence**

→ Evidential Uncertainty

[1] Alex Kendall and Yarin Gal. **What Uncertainties Do We Need in Bayesian Deep Learning for Computer Vision?** NIPS 2017.
[2] Audun Jøsang, Jin-Hee Cho, and Feng Chen. **Uncertainty Characteristics of Subjective Opinions.** FUSION 2018.

Evidential Uncertainty

Task: 3 class image classification

Training Data: Dog (e_1 =10 images) Cat (e_2 = 10 images) Pig (e_3 = 10 images)

$e = [e_1, \ldots, e_K]$ Evidence (Historical observations)

$b_i = \frac{e_i}{\sum_{k=1}^{K} \alpha_k}$ Belief

$O = [b_1, \cdots, b_K, u]$ Opinion

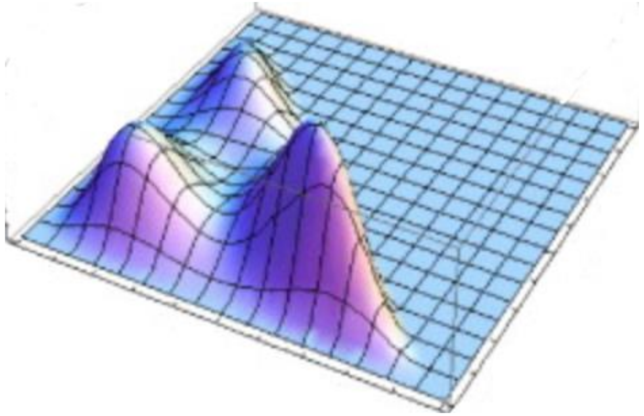

where u is vacuity and α is the strength of each singleton belief.

$\alpha = [\alpha_1, \cdots, \alpha_K]; \alpha_k = r_k + a_k \cdot W$

$\alpha = e + 1$ with $W = K$ and $a_k = 1/K$

A subjective opinion modeled based on `Subjective Logic' which uses Dirichlet distribution to measure multiple dimensions of uncertainty in classification tasks

Why Evidential Uncertainty?

	Confidence Prediction	Conflict Prediction	Out-of-Distribution
	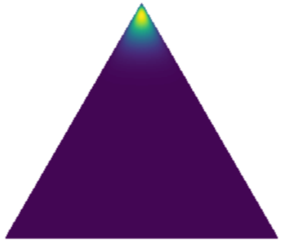	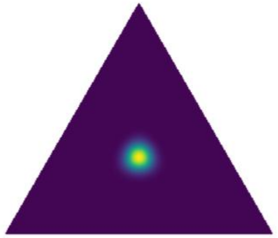	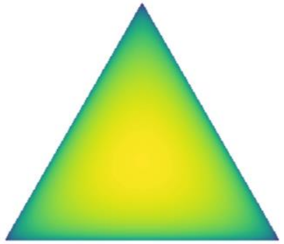
Dirichlet Distribution	α = [11, 1, 1]	α = [11, 11, 11]	α = [1, 1, 1]
Expected Probability	p = [0.83, 0.083, 0.083]	p = [1/3, 1/3, 1/3]	p = [1/3, 1/3, 1/3]
	Low Uncertainty	High Dissonance (conflicting evidence)	High Vacuity (lack of evidence)
Test image			

Different Vacuity

Sample probability

How to train a model estimating the Dirichlet distribution with evidential uncertainty?

Evidential Neural Networks

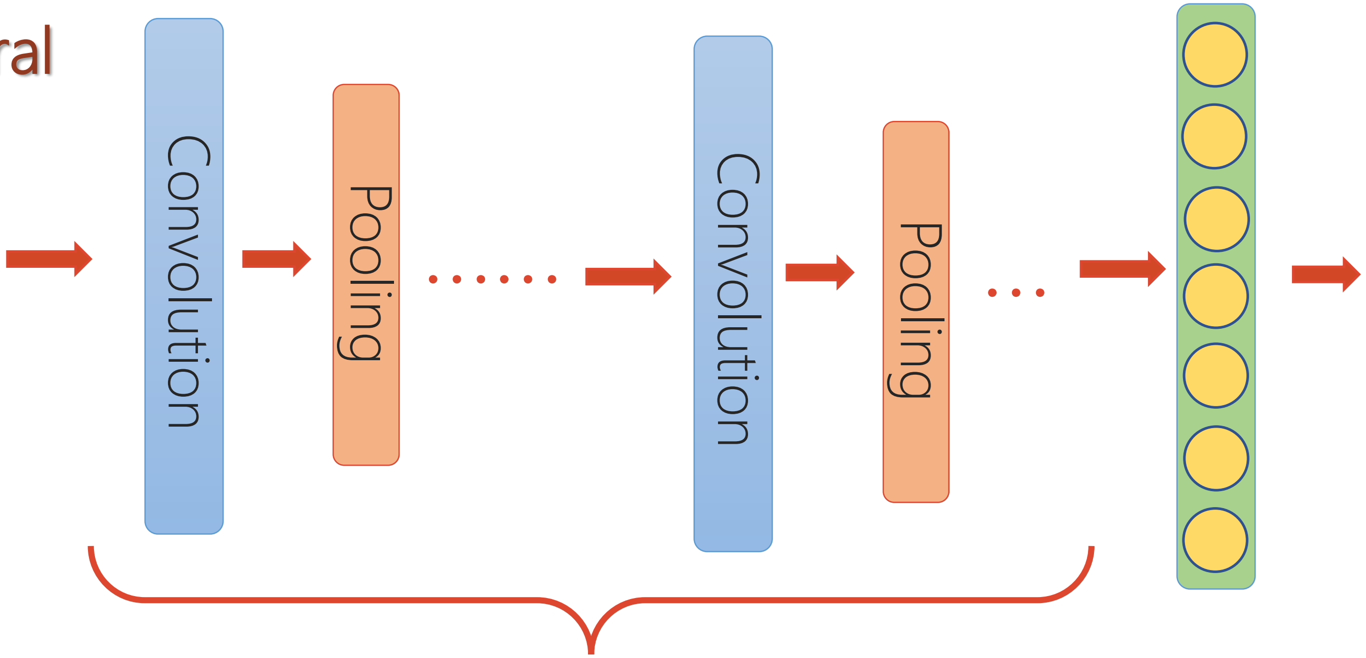

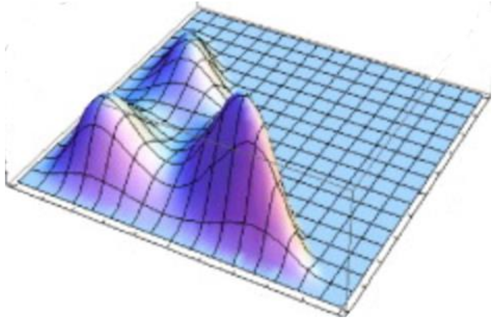

Dirichlet Distribution

CNN model

$f(\mathbf{x}_i, \Theta) = \alpha$

Higher classification accuracy and Uncertainty estimation

Square Loss:

$$\mathcal{L}(f(\mathbf{x}_i|\Theta), \mathbf{y}_i) = \int \|\mathbf{y}_i - \mathbf{p}_i\|_2^2 \mathrm{Prob}\Big(\mathbf{p}_i | f(\mathbf{x}_i, \Theta)\Big) d\mathbf{p}_i = \sum_{j=1}^{K} \big(y_{ij} - \mathbb{E}[p_{ij}]\big)^2 + \mathrm{Var}(p_{ij})$$

Expected Probability

Prediction error Variance

Minimize

Regularized Evidential Neural Networks for Quantifying Uncertainty

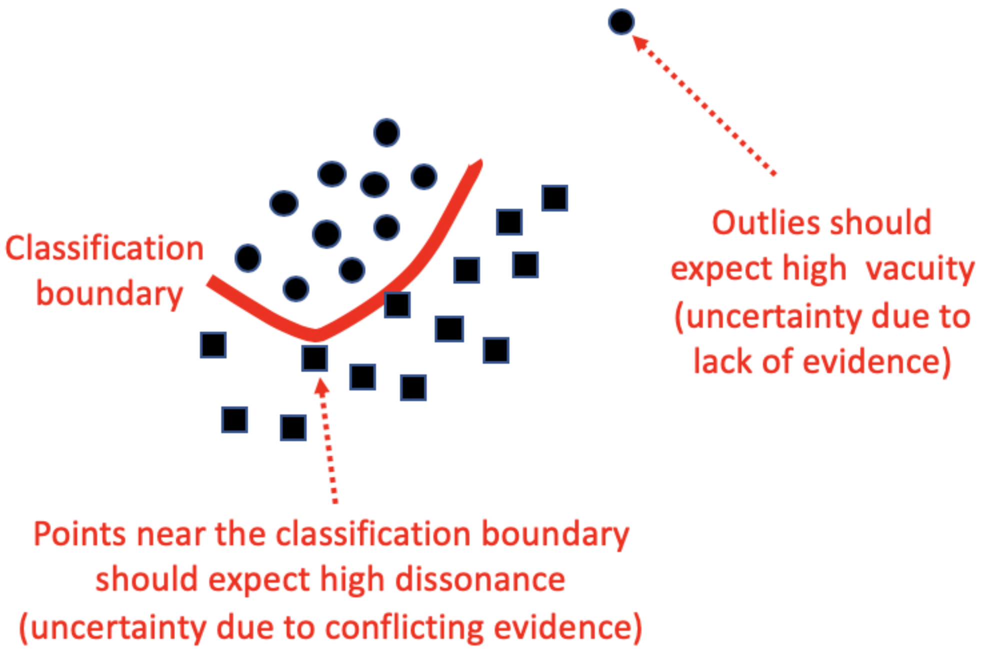

Out-of-Distribution (OOD) samples	Boundary samples
↓	↓
High Vacuity	High Dissonance

$$\mathcal{L}(\Theta) = \mathbb{E}_{(\mathbf{x}_i,\mathbf{y}_i)\sim\mathcal{D}}[\mathcal{L}(f(\mathbf{x}_i|\Theta),\mathbf{y}_i)] - \lambda_1 \mathbb{E}_{(\mathbf{x}_i,\mathbf{y}_i)\sim\mathcal{D}_{\text{OOD}}}[\text{Vac}(f(\mathbf{x}_i|\Theta))] - \lambda_2 \mathbb{E}_{(\mathbf{x}_i,\mathbf{y}_i)\sim\mathcal{D}_{\text{BOD}}}[\text{Diss}(f(\mathbf{x}_i|\Theta))]$$

- Define the outlier samples for testing OOD
- Minimize the loss function by maximizing vacuity for OOD and dissonance for boundary samples, in addition to the normal loss function to maximize prediction accuracy.

Quantifying vacuity and dissonance under different ENN models

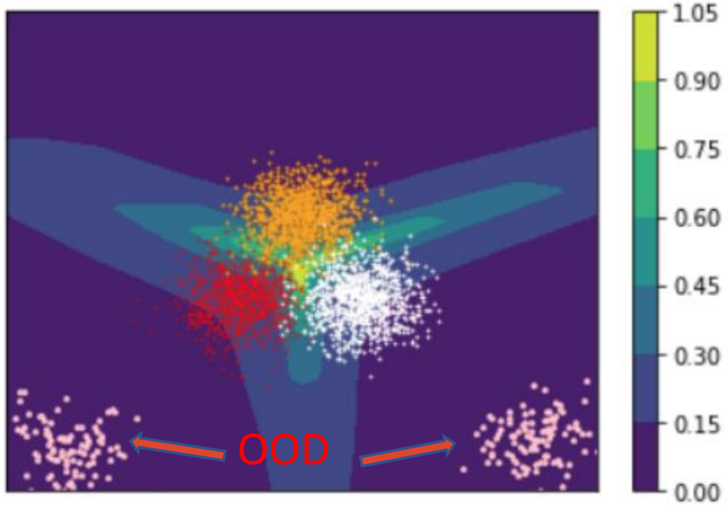

(a) Vacuity contour (ENN model)

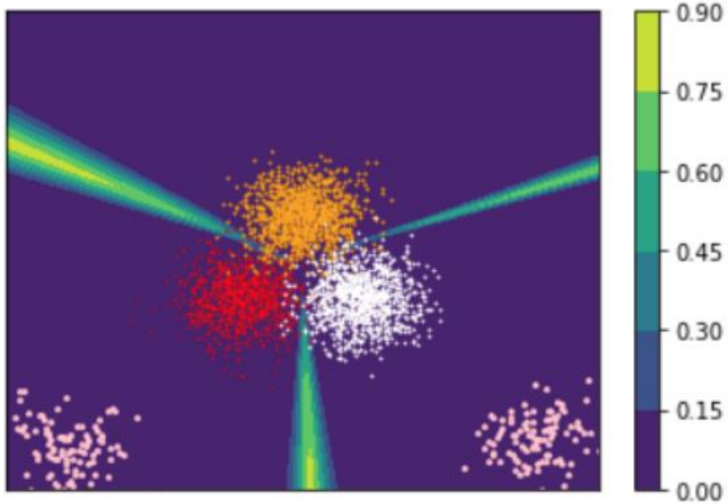

(b) Dissonance contour map (ENN model)

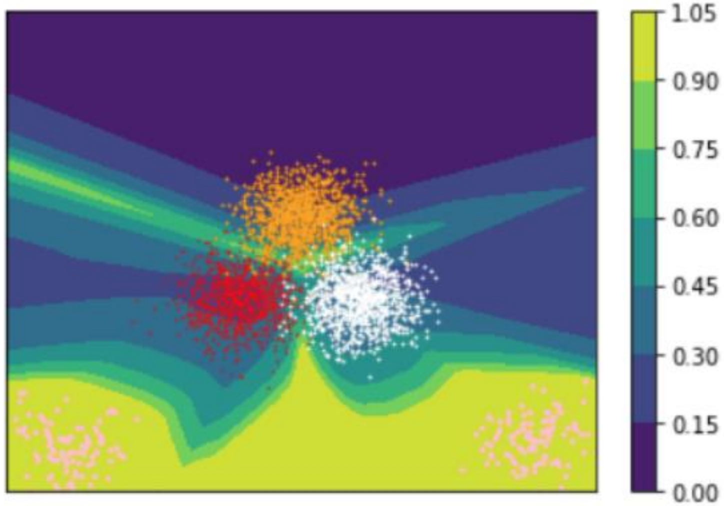

(c) Vacuity contour map (ENN-Vac model)

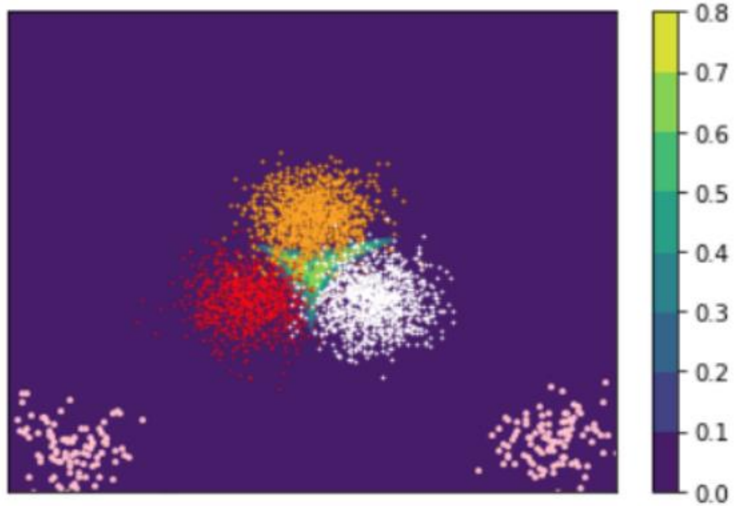

(d) Dissonance contour map (ENN-Vac-Diss model)

Setting:

- Generated a synthetic dataset with three classes: red, orange, and white.
- Under a given NN model, we sampled 1000 points of each class based on the Gaussian distribution in a 2-D space.
- We used 200 total OOD samples to train regularized-ENN models.

Results:

- Under regularized-ENN models, **high vacuity in OOD region and high dissonance on in-class boundary** are observed.

OOD Detection vs. Entropy

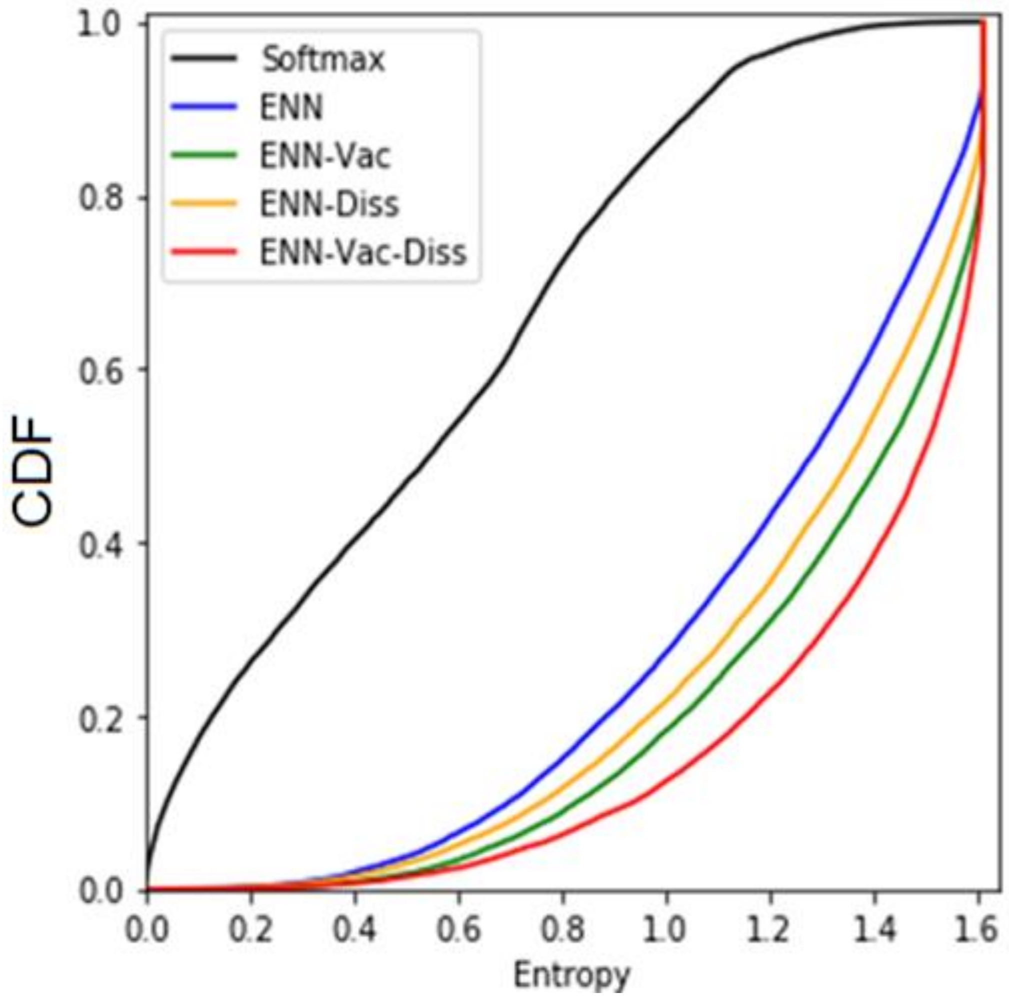

Setting:

10 classes of Cifar10 dataset are divided in 3 groups:

{airplane, automobile, bird, cat, deer}: training and validation

{ship, truck}: OOD training

{dog, frog, horse}: OOD detection test

Result:

- Right most corner is more desirable in OOD (higher entropy is associated with higher OOD detection)
- Higher vacuity (e.g., ENN-Vac with green; ENN-Vac-Diss with red) is obviously related to higher OOD.
- The proposed **regularized-ENN models show clearer uncertainty effect in OOD detection** than the baseline models.

Conclusions

- We **proposed regularized Evidential Neural Networks (ENN)** considering evidential uncertainty, **vacuity** and **dissonance**.
- We showed **anticipated predictive measurements of vacuity and dissonance in our proposed ENN-Vac and ENN in out-of-distribution (OOD) and boundary samples**.
- We validated the **outperformance of our proposed method (i.e., ENN-Vac-Diss)** over other schemes in terms of the performance in the **OOD detection** task.
- We **bridged a belief/evidence model with deep learning to predict multidimensional uncertainty.**

Open Research Questions

- Why is predicting uncertainty important?
- What is knowing an extent of uncertainty useful for decision making?
- What types of uncertainty are more important than others?
    - Is lack of information better than wrong information?
    - Is perceiving high uncertainty better than misclassification?

Thank you!

Any Question & Comments?

Jin-Hee Cho, CS@VT
jicho@vt.edu