A Rate-Distortion View of Uncertainty Quantification

Ifigeneia Apostolopoulou, Benjamin Eysenbach, Frank Nielsen, Artur Dubrawski

Code and paper!

https://github.com/ifiaposto/Distance_Aware_Bottleneck

Problem: Efficient Uncertainty Quantification

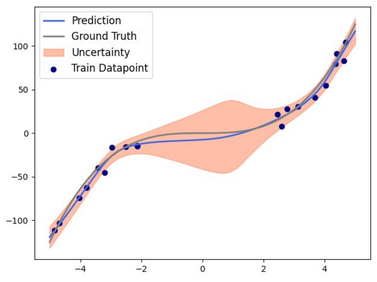

Goal: reliable & lightweight models that "know what they know."

- **Principled**: How much evidence?
- **Efficient**: Single Model & Sampling-Free
- **Post-hoc**: Works on pre-trained models

Key Idea: Distance-Awareness

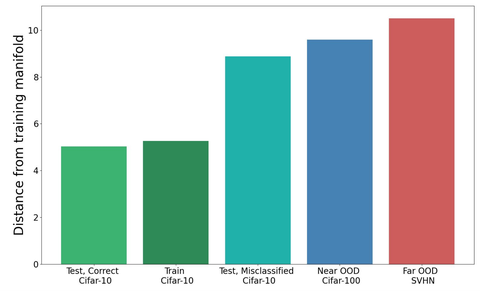

How far is an input example from the training dataset?

Our Method: Distance-Aware Bottleneck (DAB)

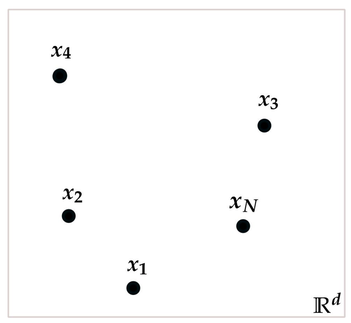

1> Treat the training dataset as a source generating training examples.

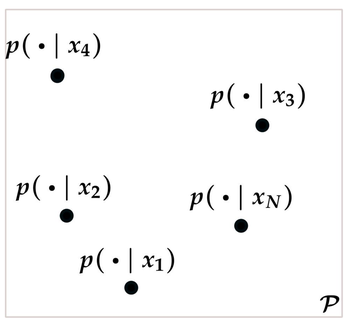

2> Map each datapoint to a *distribution* through a stochastic encoder.

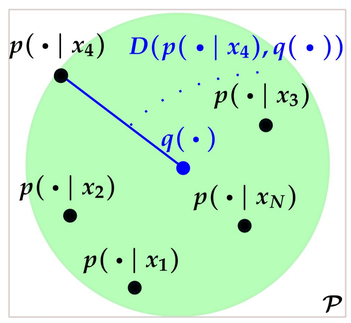

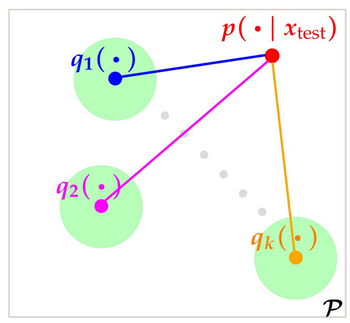

3> The codebook contains the centroids of the encoders in terms of a *statistical distance*. Uncertainty is quantified by the expected distance from the codebook: $\text{uncertainty}(x_{\text{test}}) = \mathbb{E}[D(p(\boldsymbol{z} \mid \boldsymbol{x}_{\text{test}}; \boldsymbol{\theta}), q_\kappa(\boldsymbol{z}; \boldsymbol{\phi}))]$

Codebook Visualization

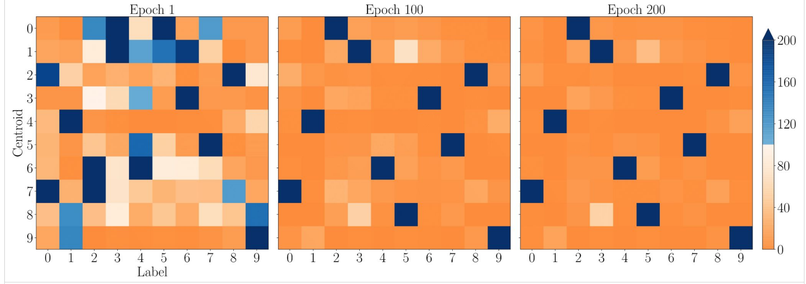

Number of CIFAR-10 test datapoints assigned to each centroid during training. Each centroid progressively attracts datapoints of the same class.

DAB's performance on ImageNet 1-K

Method	Uncertainty Description	Calibration AUROC ↑	OOD AUROC ImageNet-O ↑	Accuracy ↑	# Trainable Parameters
Deep Ensemble of 5	Gibbs softmax entropy	0.861 ± 0.0004	0.642 ± 0.001	**78.4 ± 0.06%**	117,672,960
DAB with fine-tuned ResNet-50 (ours)	Statistical distance (KL)	$\mathbf{0.868 \pm 0.0008}$	$\mathbf{0.743 \pm 0.004}$	$76.1 \pm 0.02\%$	36,612,328
DAB with pre-trained ResNet-50 (ours)	Statistical distance (KL)	0.866 ± 0.0003	0.732 ± 0.004	$74.71 \pm 0.09\%$	**13,077,736**

DAB outperforms ensembles at predicting misclassifications!

DAB can better distinguish out-of-distribution images.

Many fewer parameters!

References

[1] Alemi, A. A., Fischer, I., Dillon, J. V., and Murphy, K. Deep variational information bottleneck. In International Conference on Learning Representations, 2017.
[2] Alemi, A. A., Fischer, I., and Dillon, J. V. Uncertainty in the variational information bottleneck. arXiv preprint arXiv:1807.00906, 2018
[3] Liu, J., Lin, Z., Padhy, S., Tran, D., Bedrax Weiss, T., and Lakshminarayanan, B. Simple and principled uncertainty estimation with deterministic deep learning via distance awareness. Advances in Neural Information Processing Systems, 2020.
[4] Banerjee, A., Merugu, S., Dhillon, I. S., Ghosh, J., and Lafferty, J. Clustering with Bregman divergences. Journal of Machine Learning Research, 2005.
[5] Van Amersfoort, J., Smith, L., Teh, Y. W., and Gal, Y. Uncertainty estimation using a single deep deterministic neural network. In International Conference on Machine Learning, 2020.