



Yale

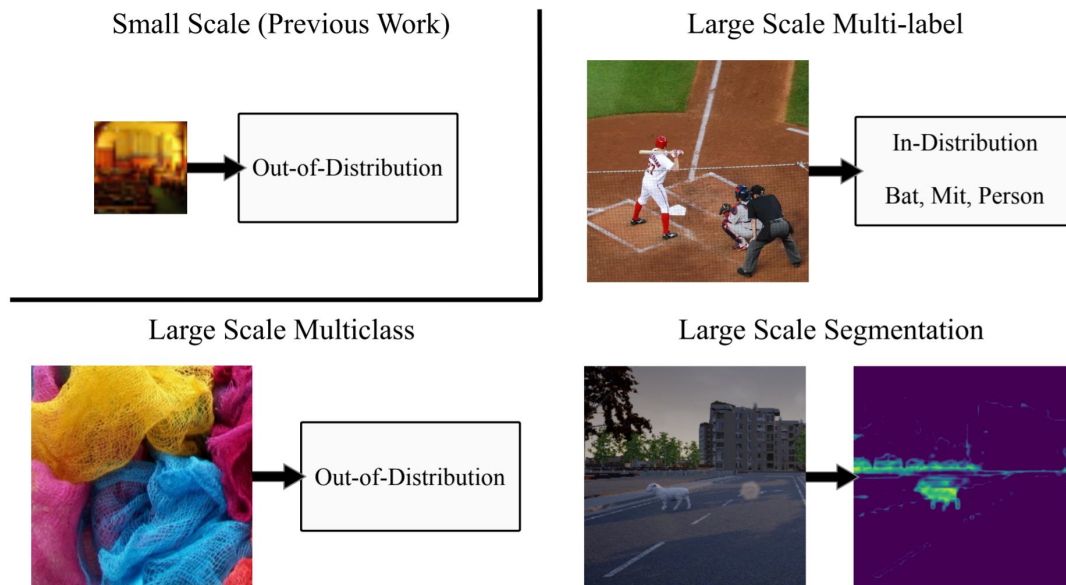


Scaling Out-of-Distribution Detection for Real-World Settings

Dan Hendrycks*, Steven Basart*, Mantas Mazeika, Andy Zou,
Joe Kwon, Mohammadreza Mostajabi, Jacob Steinhardt, Dawn Song

Overview



- We investigate out-of-distribution (OOD) detection in large-scale settings
- We identify several problems and solutions specific to large-scale settings
- We introduce new datasets and a new baseline for large-scale OOD detection



A problem with MSP

A problem with MSP

- Scaling up to datasets with many fine-grained classes dilutes probability mass
- Solution: Instead of maximum softmax probability (MSP), use MaxLogit



MaxLogit results

When trained on ImageNet or Places365 (\mathcal{D}_{in}) MaxLogit outperforms MSP (averaged across various OOD test sets ($\mathcal{D}_{\text{out}}^{\text{test}}$))

\mathcal{D}_{in}	FPR95 ↓			AUROC ↑			AUPR ↑		
	MSP	DeVries	MaxLogit	MSP	DeVries	MaxLogit	MSP	DeVries	MaxLogit
ImageNet	44.2	46.0	35.8	84.6	76.9	87.2	38.2	30.5	45.8
Places365	52.6	85.8	36.6	76.0	31.1	85.8	8.2	2.0	19.2

OOD detection with ImageNet-21K models

Problem: ImageNet-21K overlaps with many datasets

- Pretraining on ImageNet-21K and using CIFAR-10/100 as $\mathcal{D}_{\text{out}}^{\text{test}}$ is invalid, as these datasets have semantic overlap
- Several prior works finds that Vision Transformers pretrained on ImageNet-21K are highly effective at OOD detection...

...but these works use CIFAR-10/100 as $\mathcal{D}_{\text{out}}^{\text{test}}$!

- 1) Is this a problem?
- 2) How can we fix it?

The Species out-of-distribution dataset

- Hundreds of anomalous species
- Disjoint with ImageNet-21K



Species results


- All models are pretrained on ImageNet-21K
- Vision Transformers (ViT) are not especially effective
- But MaxLogit reliably improves AUROC!

\mathcal{D}_{in}	\mathcal{D}_{out}^{test}	ResNet		ViT		MLP Mixer	
		MSP	MaxLogit	MSP	MaxLogit	MSP	MaxLogit
ImageNet-21K-P	Amphibians	40.1	48.3	41.3	49.0	42.7	50.1
	Arachnids	45.6	54.6	44.8	55.0	47.1	57.2
	Fish	40.6	55.5	41.2	53.6	41.8	53.4
	Fungi	66.0	76.8	63.9	76.1	63.7	76.4
	Insects	46.8	54.9	47.6	52.8	47.8	52.1
	Mammals	45.0	50.0	47.6	47.5	48.1	46.3
	Microorganisms	76.3	82.4	69.3	81.0	72.7	84.9
	Mollusks	44.5	51.9	43.4	49.8	44.8	51.6
	Plants	68.4	75.8	65.7	72.9	67.2	73.9
	Protozoa	72.9	81.6	71.8	81.8	71.2	79.1
Mean		54.6	63.2	53.7	61.9	54.7	62.5

Anomaly segmentation


StreetHazards

- Hundreds of unique anomalous objects
- Naturally integrated into driving scenes



StreetHazards results

- MaxLogit outperforms numerous baselines quantitatively
- Compared to MSP, MaxLogit reduces edge artifacts
- Room for improvement, especially in AUPR



		MSP	MaxLogit
FPR95	↓	33.7	26.5
AUROC	↑	87.7	89.3
AUPR	↑	6.6	10.6

Summary

- Large-scale settings present new challenges for OOD detection
- We introduce new datasets for large-scale settings (Species and StreetHazards)
- We identify a simple baseline that performs especially well in large-scale settings (MaxLogit)
- For more experiments, including multi-label, please see our paper!

Thank you