Classifying dogs using PAWS

Semi-supervised fine-grain classification

Umesh Thillaivasan — June 4, 2022 — CS231n Spring 2022

Stanford University

Introduction | Learning with sparsely labeled data

Challenge: fine-grain classification with sparsely labeled datasets

Given: a large set of unlabeled data, and a small set of labeled data

Goal: learn good visual representations that generalize to unseen data

Semi-supervised learning by using labeled data

1. Pre-train with both labeled and unlabeled data
2. Fine-tune with labeled data

Datasets | CIFAR-10 and ImageNet

CIFAR-10

10 classes
60K images

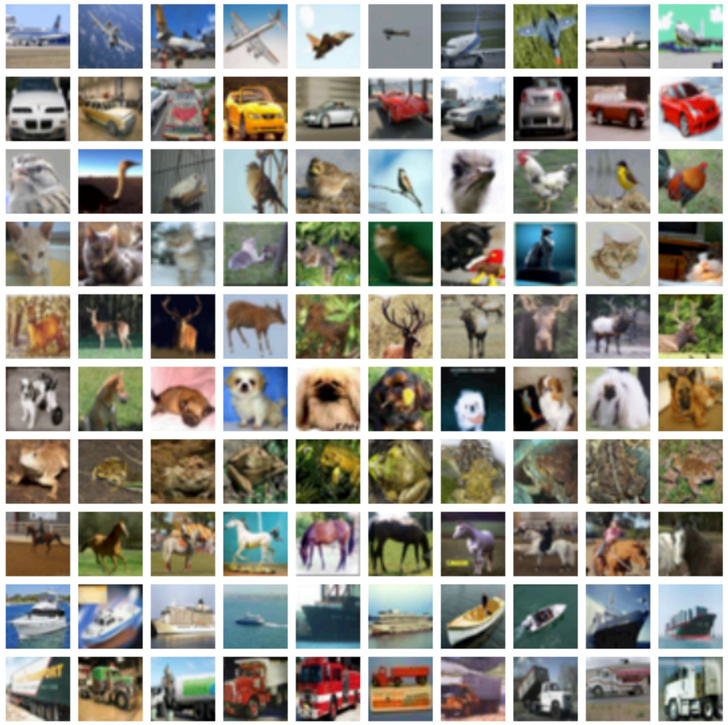

ImageNet

1000 classes
1.2M images

Stanford Dogs

120 classes
20K images

Figures of CIFAR-10, ImageNet, and Stanford Dogs datasets

Method | Predicting view Assignments With Support samples

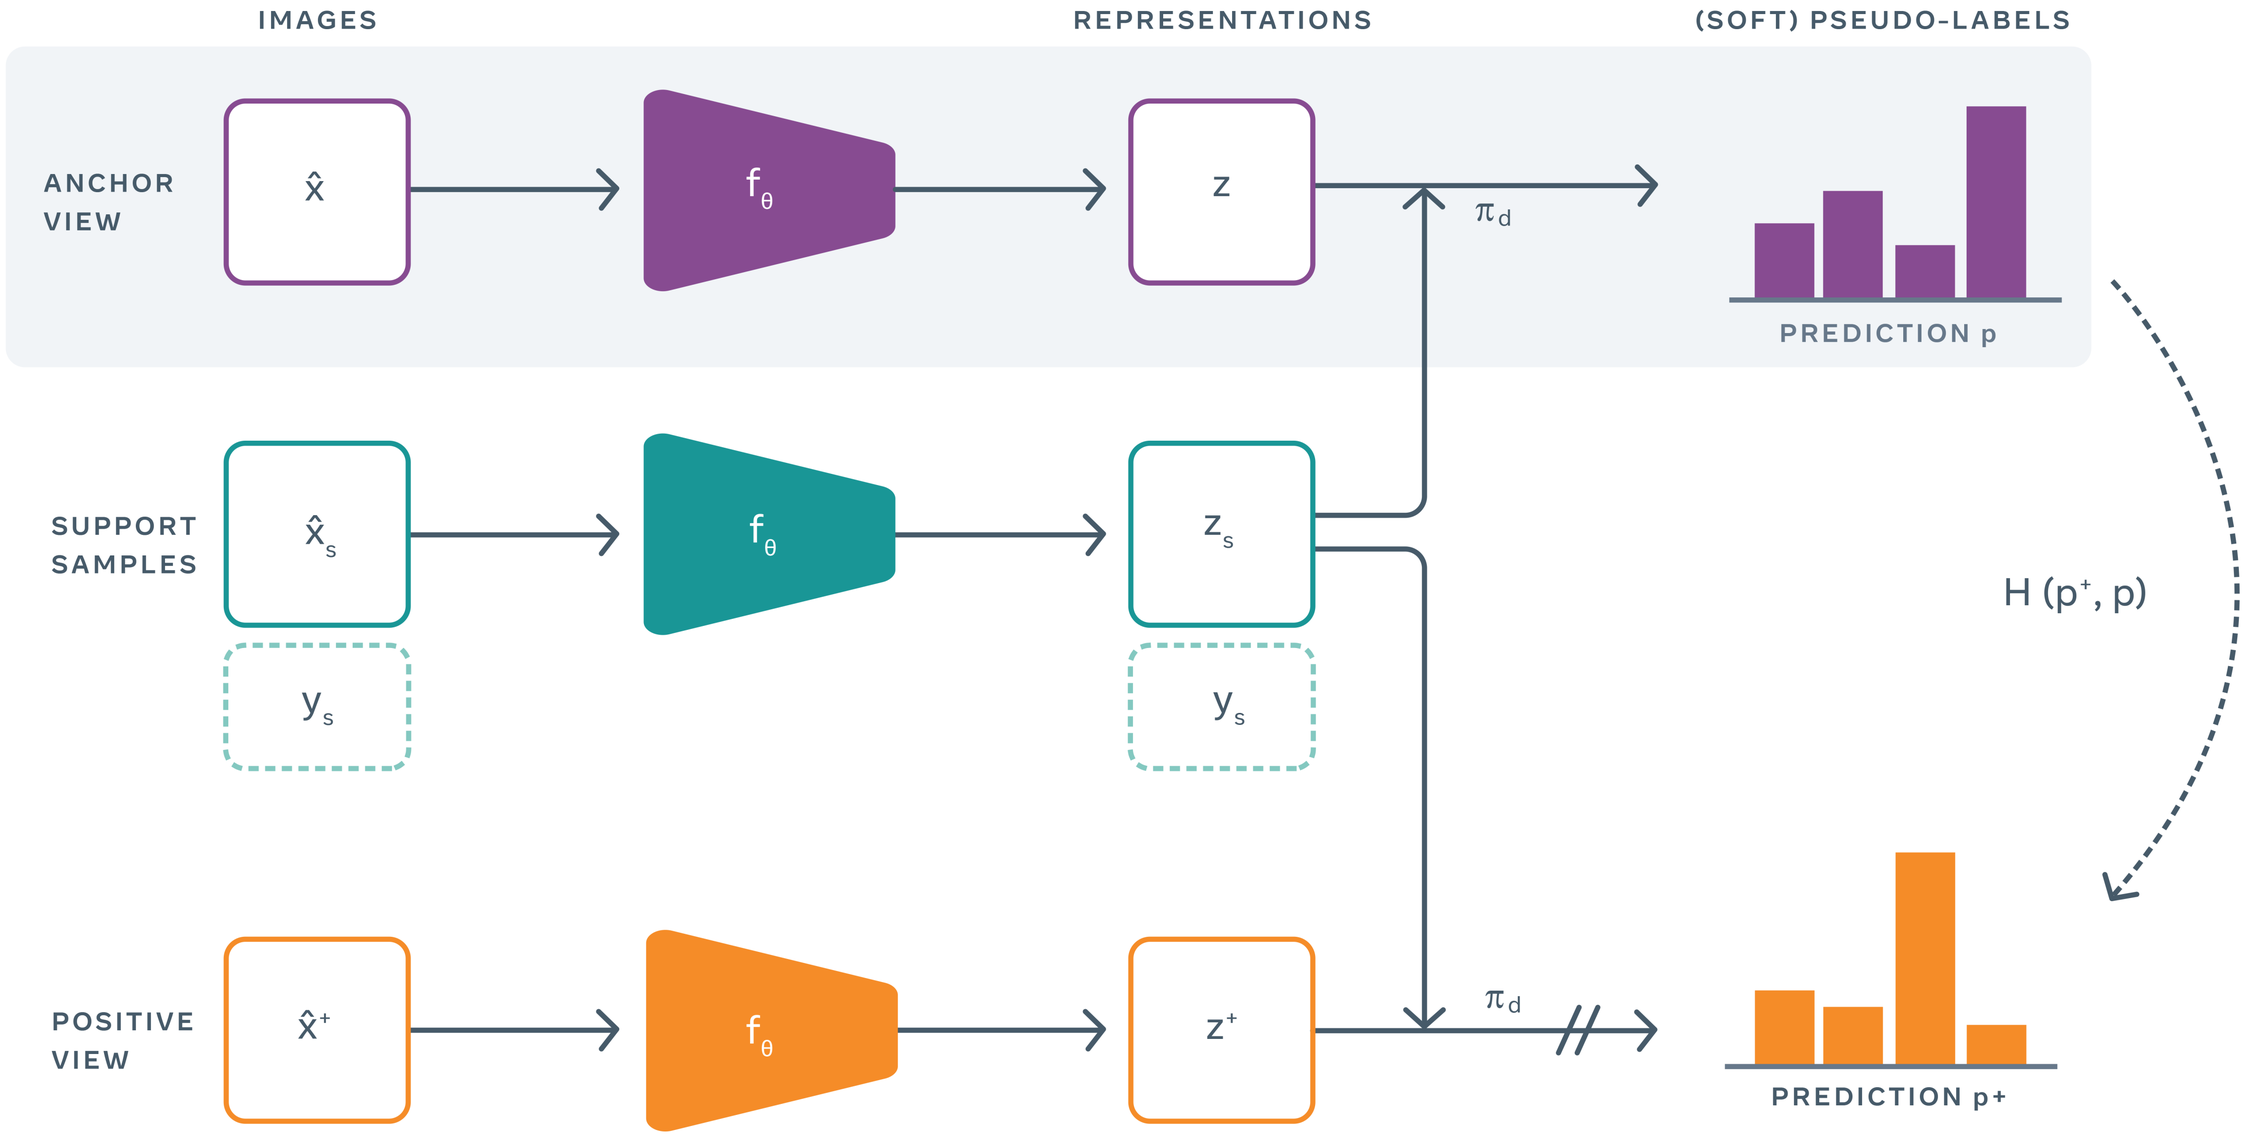

Figure of PAWS architecture from PAWS Github repository, Assran et al

Experiments | Fine-tuning while amount of labeled data

PAWS SNN accuracy did not significantly drop as sparsity became very large

Experiments	Fine-tuning				Soft Nearest Neighbours		
Labeled : Unlabeled Ratio	Epochs	Training Accuracy	Validation Accuracy	Delta to Best Val Acc.	Top 1 Accuracy	Top 5 Accuracy	Delta to Best Top 1 Acc
1:1 Ratio	30	87.56%	83.30%	0.00%	59.42%	87.13%	0%
1:2 Ratio	30	87.51%	81.89%	-1.41%	56.45%	85.72%	-2.97%
1:5 Ratio	30	87.40%	80.96%	-2.34%	53.7%	83.73%	-5.72%
1:10 Ratio	30	87.34%	80.86%	-2.44%	52.52%	83.08%	-6.9%
1:25 Ratio	30	86.80%	81.20%	-2.10%	52.18%	81.63%	-7.24%

Analysis | Saliency Maps

Learned representations of fine-grain classes with high sparsity

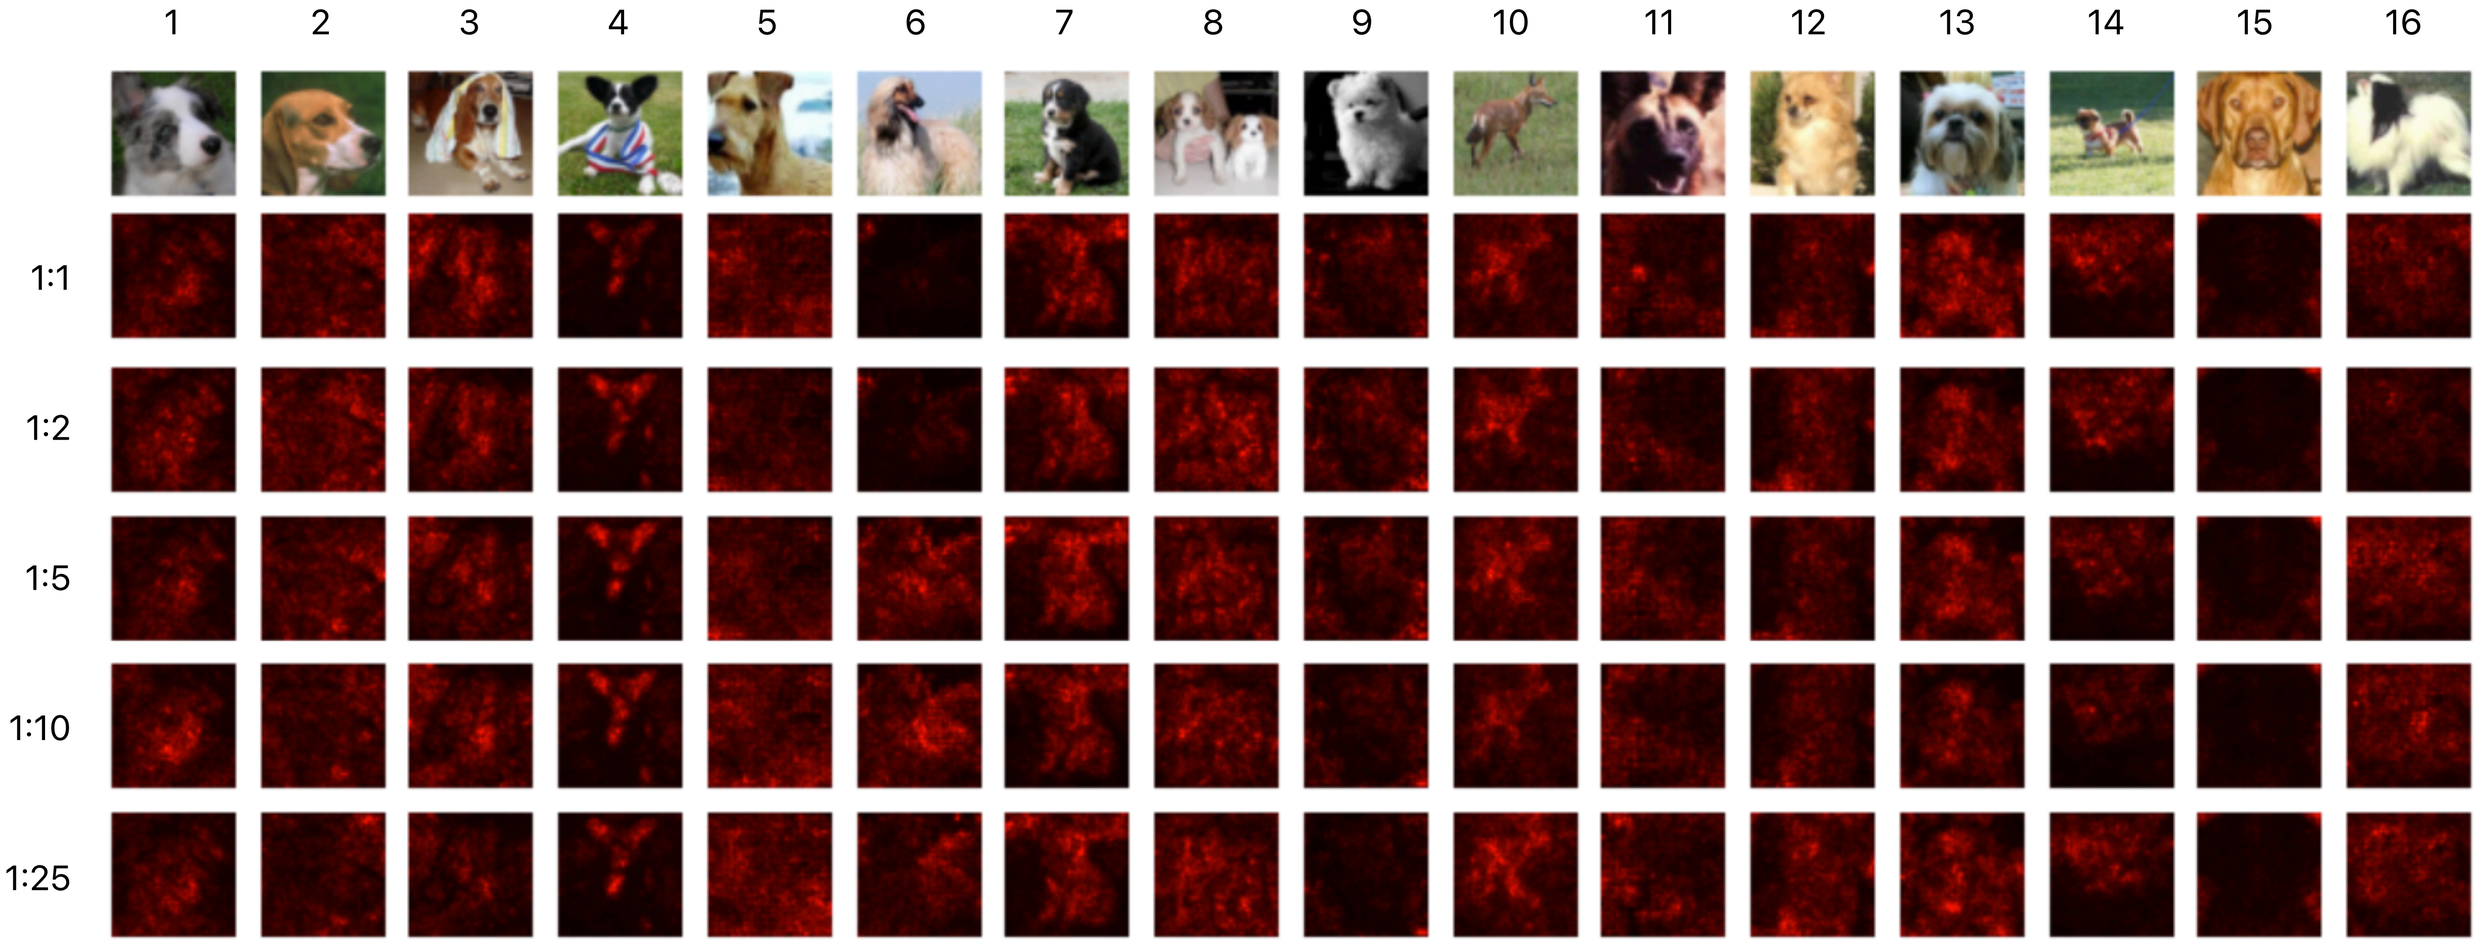

Analysis | Confusion Matrices

As sparsity increased, class 3 and 89 generated significant false predictions

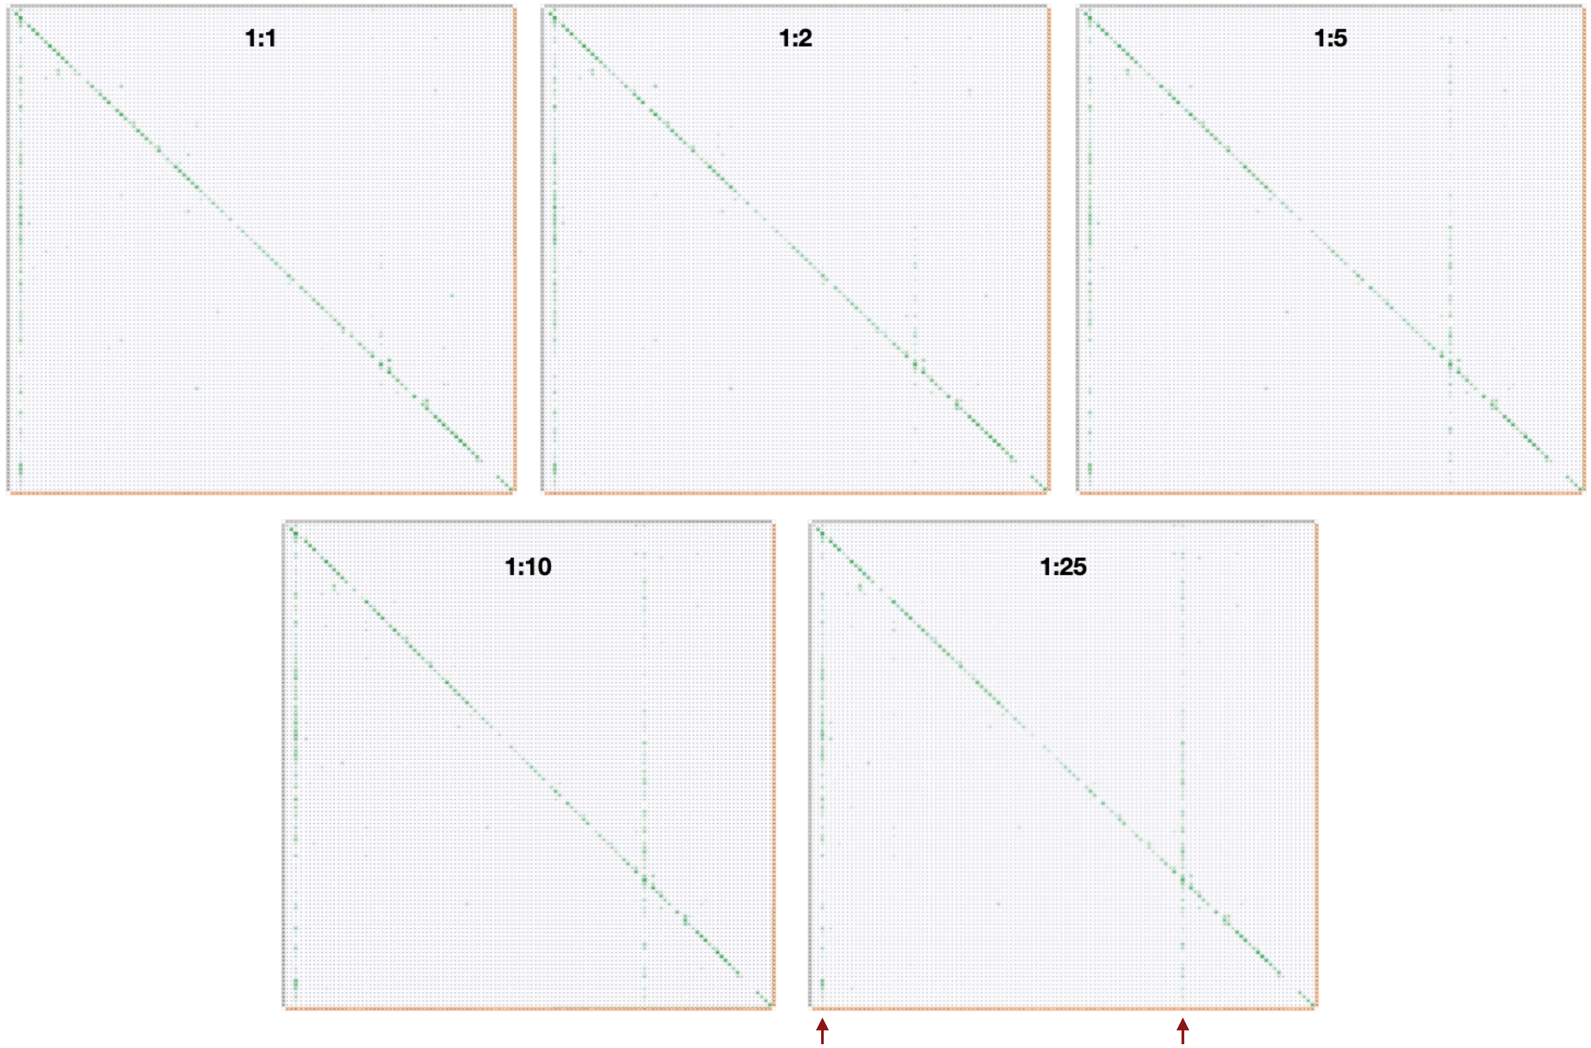

Class 3 — Maltese

Class 89 — Bernese Mountain Dog

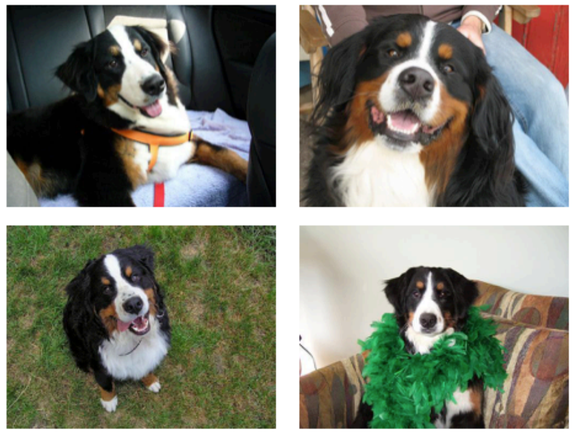

Conclusion | PAWS is viable for fine-grain classification

PAWS can use sparsely labeled data to effectively learn a new fine-grain class

Performance diminished when using very sparse labels (>1:5)

Future work is to fine-tune on other fine-grain datasets

Classifying dogs using PAWS

Semi-supervised fine-grain classification

Umesh Thillaivasan — June 4, 2022 — CS231n Spring 2022

Stanford University