Few-Shot Image Classification Along Sparse Graphs
Supplementary Material

Joseph F Comer Philip L Jacobson
Heiko Hoffmann
HRL Laboratories, LLC
3011 Malibu Canyon Rd, Malibu, CA 90265
joseph.comer@ncf.edu, philip_jacobson@berkeley.edu, drawfind@gmail.com

In this supplementary material, we present 1) the dependence of K-Prop on the number of propagated labels for 1-shot learning, 2) the relationship between the 1-shot accuracy and the probability that nearest neighbors are in the same class, 3) additional comparisons between K-Prop and the three other methods from the literature: MatchingNets, ProtoNets, and Adaptive Subspaces, and 4) numerical values for our ablation studies.

1. Dependence on number of extra labels

Figure 1 shows the dependence of the 1-shot accuracy for K-Prop (with EsViT backbone) on the number, M, of labels added by our label propagation method. For each number of labels added, we sampled $1,000$ random 5-way tasks and plotted the mean. We evaluated each value of M from 1 to 10, and subsequently $M = 15, 20, \dots, 100$. The maximum accuracy depended on the dataset used, e.g., for RESISC45, the peak was at $M = 30$, while it was at a lower value for CropDisease and a higher value for EuroSat. In our main experiments, we chose $M = 4$ for all datasets for simplicity and computational speed, but for EuroSat we could have gotten significantly better results using $M = 95$ instead of 4.

2. Accuracy vs nearest-neighbor probability

Figure 2 shows the absolute values for the relationship between 1-shot learning accuracy and the probability that nearest-neighbors are in the same class, p_{NN}. For each dataset, we evaluated three different backbone networks: Resnet18 trained on Imagenet1k, EsViT trained on Imagenet1k, and EsViT trained on the target data.

3. Comparison with other methods

We present additional results comparing K-Prop to other methods, using again ProtoNets, MatchingNets, and Subspaces for comparison. Here, we used a fixed backbone (either Resnet18 or EsViT) pretrained on Imagenet1k for each method. For each of the meta-learning algorithms, we carried out meta-training using mini-Imagenet. We then evaluated each method on the RESISC45, CropDisease, EuroSat, CUB, and Fungi datasets.

For K-Prop with EsViT, no label information from Imagenet1k or mini-Imagenet is used, while for K-Prop with Resnet18, we used the same Imagenet1k pretrained weights as with the other methods, but no additional training using mini-Imagenet. The comparisons using a Resnet18 backbone are shown in Tab. 1, while the comparisons using an EsViT backbone are shown in Tab. 2. Despite K-Prop with EsViT being at a disadvantage compared to ProtoNets and MatchingNets, it still outperformed both in most cases. For this comparison with the Imagenet-pretrained EsViT, we omitted Subspaces due to computation-time limits.

4. Ablation study

We provide the results of Fig. 9 from the main paper in tabular form in Tables 3, 4, and 5.

# of Shots	ProtoNets	MatchingNets	Subspaces	Ours
RESISC				
1	49.47 ± 0.8	40.45 ± 0.7	49.14 ± 0.8	$\mathbf{71.15 \pm 0.7}$
2	57.45 ± 0.7	53.50 ± 0.7	55.57 ± 0.7	$\mathbf{76.52 \pm 0.6}$
5	69.9 ± 0.6	64.03 ± 0.6	67.12 ± 0.6	$\mathbf{84.84 \pm 0.5}$
CropDisease				
1	57.16 ± 0.9	46.36 ± 0.9	51.01 ± 0.9	$\mathbf{78.53 \pm 0.7}$
2	69.76 ± 0.8	53.75 ± 0.9	55.28 ± 0.8	$\mathbf{83.56 \pm 0.6}$
5	80.77 ± 0.7	62.38 ± 0.9	70.37 ± 0.7	$\mathbf{91.20 \pm 0.4}$
EuroSat				
1	60.58 ± 0.9	48.04 ± 0.8	45.94 ± 0.8	$\mathbf{67.90 \pm 0.7}$
2	67.38 ± 0.8	60.66 ± 0.8	53.61 ± 0.8	$\mathbf{74.54 \pm 0.6}$
5	77.67 ± 0.7	68.50 ± 0.8	69.79 ± 0.7	$\mathbf{82.86 \pm 0.4}$
CUB				
1	38.95 ± 0.8	37.42 ± 0.8	35.60 ± 0.7	$\mathbf{81.08 \pm 0.8}$
2	46.02 ± 0.8	42.36 ± 0.7	40.24 ± 0.7	$\mathbf{83.77 \pm 0.7}$
5	56.62 ± 0.8	49.19 ± 0.8	48.92 ± 0.7	$\mathbf{89.39 \pm 0.6}$
Fungi				
1	36.19 ± 0.7	34.17 ± 0.7	31.76 ± 0.7	$\mathbf{42.87 \pm 1.1}$
2	42.25 ± 0.8	38.00 ± 0.7	33.92 ± 0.7	$\mathbf{56.93 \pm 1.2}$
5	52.05 ± 0.7	44.93 ± 0.7	41.65 ± 0.7	$\mathbf{64.21 \pm 1.2}$

Table 1: Comparing the performance of our method with ProtoNets, MatchingNets, and Adaptive Subspaces, all using Resnet18 backbones pre-trained on Imagenet1k.

# of Shots	ProtoNets	MatchingNets	Ours
RESISC			
1	72.82 ± 0.8	72.56 ± 0.9	$\mathbf{78.60 \pm 1.0}$
2	82.08 ± 0.7	78.56 ± 0.8	$\mathbf{82.80 \pm 0.9}$
5	$\mathbf{89.48 \pm 0.5}$	85.55 ± 0.6	89.28 ± 0.9
CropDisease			
1	81.97 ± 0.8	81.94 ± 0.8	$\mathbf{86.23 \pm 0.6}$
2	89.54 ± 0.6	87.37 ± 0.7	$\mathbf{90.06 \pm 0.5}$
5	94.61 ± 0.4	92.05 ± 0.6	$\mathbf{95.43 \pm 0.3}$
EuroSat			
1	65.17 ± 0.8	64.77 ± 0.9	$\mathbf{70.85 \pm 0.8}$
2	76.22 ± 0.7	70.90 ± 0.8	$\mathbf{77.16 \pm 0.6}$
5	84.34 ± 0.5	77.36 ± 0.6	$\mathbf{84.73 \pm 0.4}$
CUB			
1	71.00 ± 1.0	71.77 ± 1.0	$\mathbf{80.00 \pm 0.8}$
2	82.23 ± 0.9	77.68 ± 0.9	$\mathbf{84.04 \pm 0.6}$
5	89.47 ± 0.7	85.54 ± 0.8	$\mathbf{90.51 \pm 0.6}$
Fungi			
1	57.26 ± 1.3	$\mathbf{57.97 \pm 1.3}$	50.24 ± 1.1
2	$\mathbf{68.11 \pm 1.1}$	66.64 ± 1.2	65.98 ± 1.1
5	$\mathbf{77.78 \pm 1.1}$	76.18 ± 1.0	71.49 ± 1.2

Table 2: Comparing the performance of our method with ProtoNets and MatchingNets, all using EsViT backbones pre-trained on Imagenet1k.

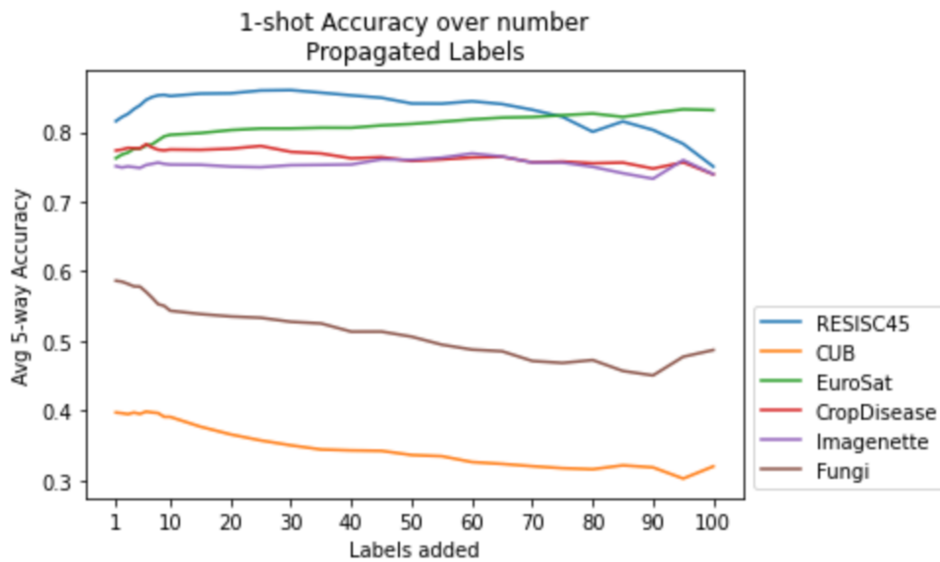

Figure 1: Dependence of 5-way, 1-shot accuracy on number of labels added using our label propagation.

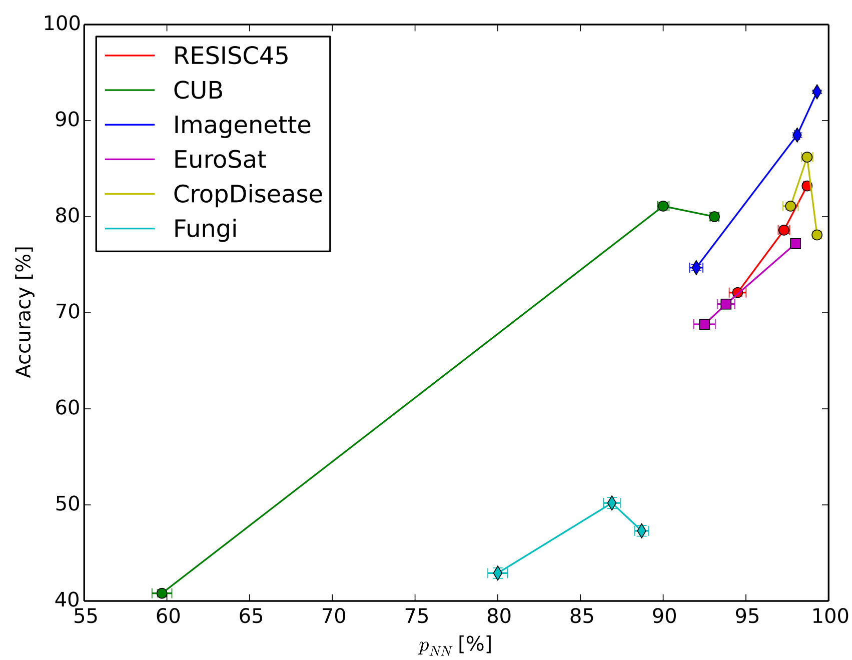

Figure 2: Accuracy for 1-shot learning vs probability that nearest-neighbors are in the same class, p_{NN} (mean $\pm$ SE). Each data point corresponds to a different dataset and backbone.

Dataset	Pretraining	Linear	LP+Linear	LP+KPCA
NWPU-RESISC45	EsViT	75.93± 0.6	80.31± 0.6	83.18± 0.6
	Pretrained ResNet18	63.29± 0.6	71.4± 0.7	72.09± 0.7
EuroSat	EsViT	69.13± 0.5	72.62± 0.5	77.2± 0.5
	Pretrained ResNet18	56.78± 0.6	66.79± 0.7	68.77± 0.7
CropDisease	EsViT	74.9± 0.8	75.77± 0.8	78.11± 0.8
	Pretrained ResNet18	72.56± 0.7	80.96± 0.7	81.12± 0.7
Fungi	EsViT	57.11± 1.0	50.61± 1.0	47.26± 1.1
	Pretrained ResNet18	48.76± 0.9	43.41± 1.0	42.87± 1.1
Imagenette	EsViT	71.78± 0.7	72.69± 0.7	74.73± 0.6
	Pretrained ResNet18	83.68± 0.4	89.74± 0.4	88.54± 0.4
CUB	EsViT	41.63± 0.6	39.79± 0.7	40.8± 0.7
	Pretrained ResNet18	68.65± 0.7	78.93± 0.8	81.08± 0.8

Table 3: 5-way, 1-shot performance of linear fine-tuning (Linear), label propagation + linear fine-tuning (LP+Linear), and label propagation + KPCA (LP+KPCA) (ours).

Dataset	Pretraining	Linear	LP+Linear	LP+KPCA
NWPU-RESISC45	EsViT	83.57± 0.5	85.12± 0.4	87.06± 0.5
	Pretrained ResNet18	73.39± 0.5	76.84± 0.6	77.1± 0.6
EuroSat	EsViT	78.01± 0.4	79.1± 0.4	83.4± 0.4
	Pretrained ResNet18	69.68± 0.5	73.52± 0.5	74.53± 0.6
CropDisease	EsViT	83.1± 0.7	83.41± 0.7	85.2± 0.7
	Pretrained ResNet18	82.81± 0.5	85.71± 0.5	86.0± 0.6
Fungi	EsViT	66.25± 0.9	64.44± 1.1	68.18± 1.4
	Pretrained ResNet18	58.24± 1.0	53.22± 1.0	56.93± 1.3
Imagenette	EsViT	79.05± 0.5	80.15± 0.5	81.12± 0.4
	Pretrained ResNet18	92.03± 0.3	93.13± 0.2	92.29± 0.3
CUB	EsViT	45.87± 0.6	45.59± 0.7	46.44± 0.9
	Pretrained ResNet18	79.01± 0.6	83.34± 0.6	83.77± 0.7

Table 4: 5-way, 2-shot performance.

Dataset	Pretraining	Linear	LP+Linear	LP+KPCA
NWPU-RESISC45	EsViT	90.6± 0.3	90.37± 0.3	92.07± 0.4
	Pretrained ResNet18	83.9± 0.4	84.45± 0.4	84.82± 0.5
EuroSat	EsViT	88.13± 0.3	87.88± 0.3	90.44± 0.3
	Pretrained ResNet18	80.96± 0.4	81.3± 0.4	82.54± 0.4
CropDisease	EsViT	91.28± 0.4	91.4± 0.4	92.82± 0.4
	Pretrained ResNet18	91.43± 0.4	91.65± 0.4	92.79± 0.4
Fungi	EsViT	77.85± 0.8	75.22± 1.0	74.29± 1.3
	Pretrained ResNet18	69.79± 0.9	63.77± 1.0	64.21± 1.3
Imagenette	EsViT	84.89± 0.3	85.12± 0.3	84.9± 0.4
	Pretrained ResNet18	96.15± 0.1	96.2± 0.1	95.97± 0.2
CUB	EsViT	54.05± 0.6	53.62± 0.6	53.59± 1.0
	Pretrained ResNet18	88.25± 0.4	88.66± 0.4	89.39± 0.7

Table 5: 5-way, 5-shot performance.