

---

# The Unreasonable Effectiveness of Fully-Connected Layers for Low-Data Regimes

## — Supplementary material —

---

**Peter Kocsis**

Technical University of Munich

[peter.kocsis@tum.de](mailto:peter.kocsis@tum.de)

**Peter Súkeník**

Technical University of Munich

[peter.sukenik@trojsten.sk](mailto:peter.sukenik@trojsten.sk)

**Guillem Brasó**

Technical University of Munich

[guillem.braso@tum.de](mailto:guillem.braso@tum.de)

**Matthias Niessner**

Technical University of Munich

[niessner@tum.de](mailto:niessner@tum.de)

**Laura Leal-Taixé**

Technical University of Munich

[leal.taixe@tum.de](mailto:leal.taixe@tum.de)

**Ismail Elezi**

Technical University of Munich

[ismail.elezi@tum.de](mailto:ismail.elezi@tum.de)

[peter-kocsis.github.io/LowDataGeneralization/](https://peter-kocsis.github.io/LowDataGeneralization/)

In this appendix, we start by providing more implementation details in Section A. In Subsection A.1, we detail the hyperparameters. In Subsection A.2, we give further information with regards to the used datasets. In Table 2, we show the number of extra parameters used during training with our method for all our experiments. In Subsection A.3, we provide the libraries used for the trainings of MLPmixer and ViT. In Subsection A.4, we provide more information for the Active Learning experiments.

In Section B, we do extra experiments with the VGG networks which already uses fully-connected layers. In Section C, we evaluate the robustness of our method on the CIFAR10-C dataset. Then, in section D, we evaluate our method on the Caltech101 dataset without any pretraining. Finally, in Section E, we provide the exact numbers (mean and standard deviation) for all the plots in the paper.

## A Training details

### A.1 Hyperparameters

Doing hyperparameter optimization in low-data regimes is extremely difficult. This is because in different cycles we have different datasets. Optimizing on them independently yields different optimal hyperparameters, which makes the results difficult to compare. Furthermore, that leads to several trainings for each cycle, which is costly. In addition, datasets like CIFAR and Caltech have public testing set, which means that an extensive hyperparameter optimization might lead to an overfitting of the testing set, and unreliable generalization results. To avoid all these issues, we decide to follow the work of LLAL [1] and use a single hyperparameter set across all the experiments and comparison. More details about the datasets and the used hyperparameters can be found in Table 1

Table 1: **Summary of the used datasets and specific hyperparameters.** The first block until the dashed line describes the datasets. The second block shows the used data augmentations, followed by the training hyperparameters. During learning rate scheduling, we divide the learning rate by 10 after 80% epochs.

	CIFAR10	CIFAR100	Caltech101	Caltech256
Number of classes	10	100	102	257
Number of training samples	50000	50000	6117	21531
Number of test samples	10000	10000	2560	9076
Official test split	✓	✓	✗	✗
Image size	32x32	32x32	Various	Various
Class balance	✓	✓	✗	✗
Random horizontal flip	✓	✓	✓	✓
Random crop	32x32, p=4	32x32, p=4	224x224, p=16	224x224, p=16
Normalization	✓	✓	✓	✓
Size of initial labeled pool	1000	5000	1000	5000
Samples labeled in stage	1000	1000	1000	1000
Number of stages	10	10	6	10
Feature size	64	64	256	256
Optimizer	SGD	SGD	SGD	SGD
Learning rate	0.1	0.1	0.001	0.001
Learning rate scheduler	✓	✓	✓	✓
Momentum	0.9	0.9	0.9	0.9
Weight decay	5e-4	5e-4	5e-4	5e-4
Number of epochs	200	200	100	100

## A.2 Datasets and the number of labels.

For all supervised experiments, we use the same pre-defined datasplits. We obtained the datasplits with incremental random sampling. This ensures that the comparisons are fair, and subject to only the network architecture. Doing otherwise might result with some method being favored by having a better training split.

## A.3 MLPMixer [2] and ViT [3].

MLPMixer [2] and ViT [3] were originally developed for ImageNet. In order to achieve their best performance we compared multiple architectures specifically tailored for the CIFAR datasets, and chose the best performing ones, specifically, MLPMixer-Nano [4] and ViT-CIFAR10 [5]

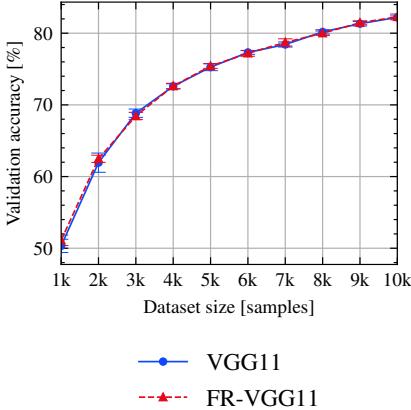
## A.4 Active Learning.

To achieve comparable results, we closely followed the active learning setup of LLAL [1]. We split the dataset into two parts: labeled and unlabeled pool. During training, we access only the labeled pool. In each cycle, we train a model on the labeled pool and use its predictions on the unlabeled pool to select samples to be included in the labeled pool for the next cycle. In each cycle, we reinitialize the model weights.

During the experiments, we noted that there is a difference between our and the officially reported LLAL [1] results. We found that the reason behind that is the mismatch between their and our PyTorch versions. They used PyTorch 1.1, while we use more modern PyTorch 1.7. Downgrading to the older version yields comparable results to the published ones. While we believe that reproducibility and comparability is crucial in research, using outdated libraries is not an ideal solution. Consequently, we compare their method and ours using PyTorch 1.7. We think that this is both fair, and allows an easier reproducibility for future research.

Table 2: **Extra parameters.** Number of extra parameters used during training for all experiments.

	CIFAR	Caltech
ResNet18	11173962	11228325
FR-ResNet18	+42058 (0.38%)	+289893 (2.58%)
ResNet34	21282122	X
FR-ResNet34	+42058 (0.20%)	X
DenseNet121	6964096	X
FR-DenseNet121	+74826 (1.07%)	X
EfficientNetB3	10711602	X
FR-EfficientNetB3	+82660 (0.77%)	X



(a) CIFAR10 plot

#Samples	VGG11	FR-VGG11
1000	50.33±0.93	<b>51.06±0.78</b>
2000	61.94±1.34	<b>62.43±0.47</b>
3000	<b>68.84±0.58</b>	68.42±0.49
4000	72.60±0.35	<b>72.62±0.35</b>
5000	75.27±0.49	<b>75.40±0.34</b>
6000	<b>77.31±0.26</b>	77.17±0.44
7000	78.42±0.36	<b>78.66±0.52</b>
8000	<b>80.18±0.31</b>	79.98±0.32
9000	81.33±0.32	<b>81.46±0.29</b>
10000	82.21±0.20	<b>82.22±0.44</b>

(b) CIFAR10 numerical results

Figure 1: **VGG experiment.** We evaluate our method with the VGG architecture. We can achieve highly similar performance with 67% fewer parameters.

## B VGG experiment

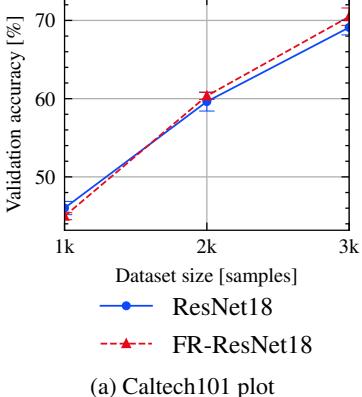
We conduct an experiment with an old-fashioned architecture that uses final fully-connected layers, the VGG11 [6]. Since this architecture follows our proposed solution, we are unable to improve upon its performance. However, we show that our light-weight Feature Refiner (FR) is able to achieve the same results, but using only 9,262,282 parameters instead of 28,144,010, resulting with a 67% reduction in the number of weights. We present the results on Figure 1.

## C Robustness benchmark

We evaluate the robustness of ResNet18 and our model (FR-ResNet18) on the CIFAR10-C dataset [7]. CIFAR10-C contains a total of 95 perturbed test sets under 19 different corruptions with 5 severity levels. We present average results of 5 independent runs (Figure 5) over the different severity levels (Figure 4) and over the different corruption types (Figure 4). Our method consistently outperforms the baseline in cases of lower corruption severity. In cases of higher corruption severity (severity 3 and severity 4), in the very low-data regime, our method significantly outperforms the baseline. However, with more added data, the baseline starts outperforming our method.

## D Caltech without pre-training

Since the Caltech datasets are more complex, and also to show that our method works with pretrained models as well, we used Imagenet-pretrained backbone for the Caltech experiments in the main paper. Now, we conduct a minimal experiment to evaluate our method without pretraining. We run our method on the Caltech101 dataset without pretraining in the first three datasplits. As it can be seen in Figure 2, in the first cycle, we are worse than ResNet18 by 1pp, similarly as with ImageNet



(a) Caltech101 plot

#Samples	ResNet18	FR-ResNet18
1000	<b>46.02</b> $\pm$ 0.85	44.98 $\pm$ 0.45
2000	59.62 $\pm$ 1.21	<b>60.35</b> $\pm$ 0.45
3000	69.10 $\pm$ 0.95	<b>70.48</b> $\pm$ 1.11

(b) Caltech101 numerical results

Figure 2: **Caltech101 w/o pretraining.** We evaluate our method on the first three splits of the Caltech101 dataset without Imagenet-pretraining. While both ResNet18 and our method reach significantly lower results than when we use ImageNet pretraining, the relative improvement of our method compared to ResNet18 remains.

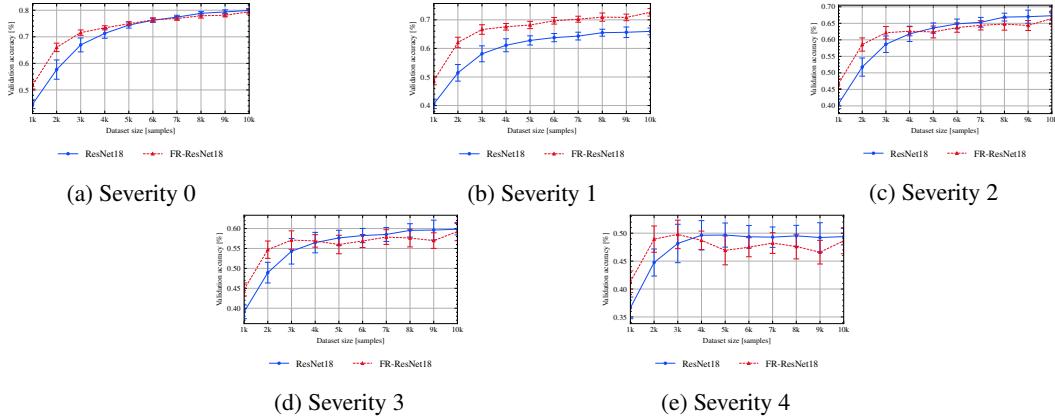


Figure 3: **Robustness benchmark - Severity.** We evaluate the robustness of our method on the CIFAR10-C dataset. We present the average score for all severity levels. Our method consistently outperforms the baseline in cases of lower corruption severity. In cases of higher corruption severity (3d and 3e), in the very low-data regime, our method significantly outperforms the baseline. However, with more added data, the baseline starts outperforming our method.

pretraining, where the difference was 0.9pp. In the second cycle, we outperform ResNet18 by 0.7pp, while with ImageNet pretraining we outperformed them by 0.5pp. In the third cycle, we outperform ResNet18 by 1.4pp, while with ImageNet pretraining we outperformed them by 0.8pp. In this way, we show that while both ResNet18 and our method reach significantly lower results than when we use ImageNet pretraining, the relative improvement of our method compared to ResNet18 remains.

## E Numerical results

In this section, we provide the exact numbers of all our results presented in the paper. Table 3, 4, 5, 6, 7, 8, 9 summarize the results provided in Figure 3, 4, 5, 6, 7, 8, 9. We provide the mean and standard deviation obtained by running the same experiment five times with different model initializations.

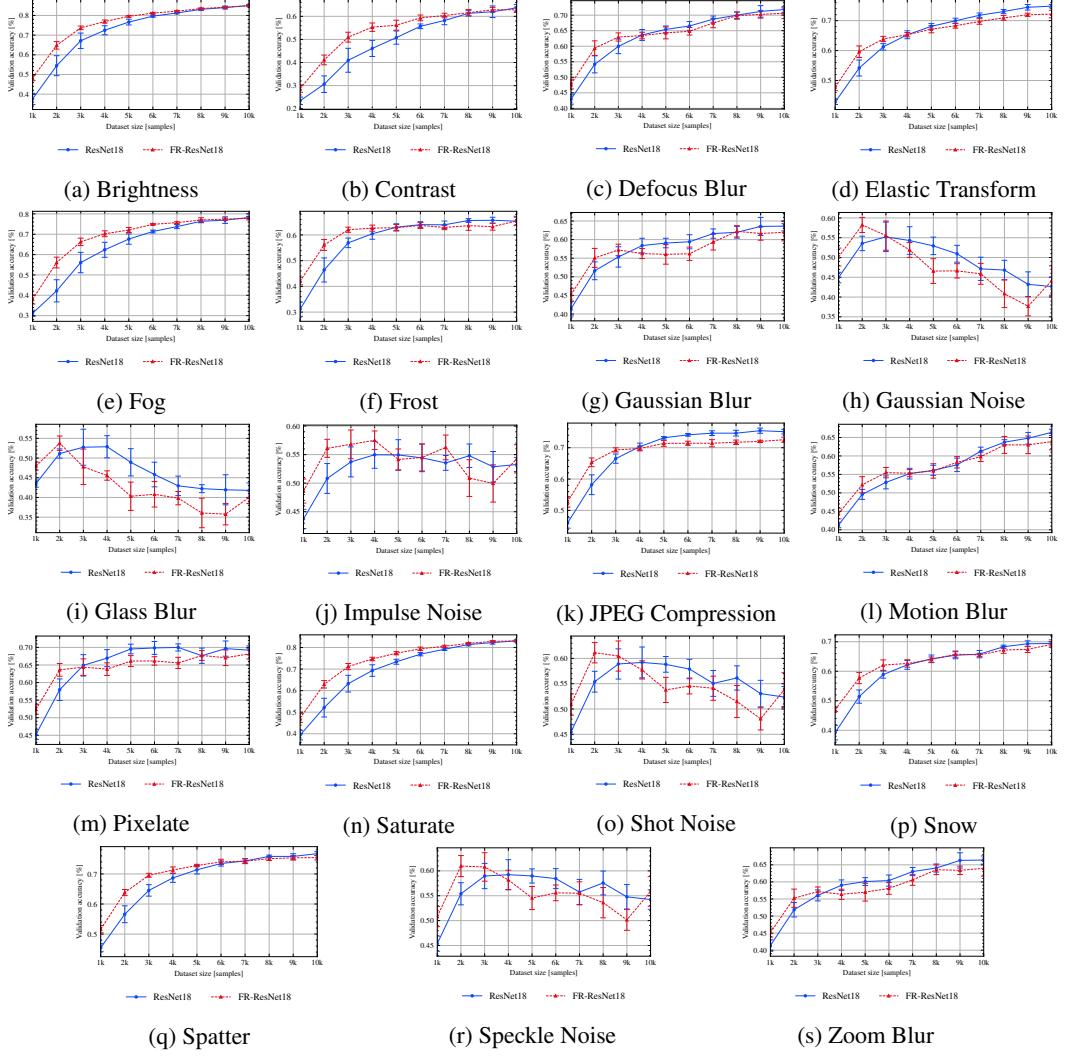


Figure 4: **Robustness benchmark - Corruption types.** We evaluate the robustness of our method on the CIFAR10-C dataset. We present the average score for all corruption types. Our method consistently outperforms the baseline in the earlier iterations.

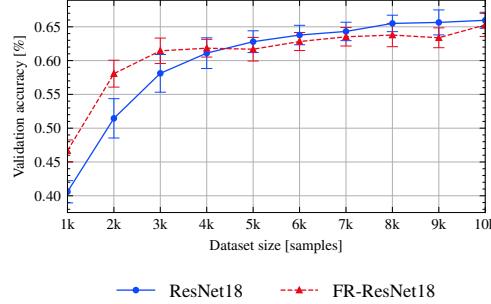


Figure 5: **Robustness benchmark - Overall.** We evaluate the robustness of our method on the CIFAR10-C dataset. We present the average score over all datasets. Our method performs better in the earlier iterations.

Table 3: **Comparisons with ResNet18 [8].** We compare our approach to our main baseline network, to ResNet18 [8] on the CIFAR10, CIFAR100, Caltech101 and Caltech256 datasets in supervised learning. Our method significantly outperforms the baseline.

(a) CIFAR10			(b) CIFAR100		
#Samples	ResNet18 [8]	FR-ResNet18	#Samples	ResNet18 [8]	FR-ResNet18
1000	46.18±2.81	<b>53.70</b> ±1.18	5000	38.07±0.92	<b>43.73</b> ±1.36
2000	60.36±1.64	<b>70.46</b> ±0.59	6000	41.72±0.61	<b>48.52</b> ±0.78
3000	69.22±3.21	<b>76.81</b> ±0.50	7000	46.29±2.07	<b>51.44</b> ±0.45
4000	75.45±2.29	<b>79.66</b> ±0.21	8000	50.06±0.60	<b>54.51</b> ±0.40
5000	80.23±0.85	<b>82.27</b> ±0.33	9000	53.25±0.32	<b>56.06</b> ±0.40
6000	82.52±0.59	<b>83.65</b> ±0.34	10000	55.84±0.90	<b>57.50</b> ±0.56
7000	84.08±0.38	<b>84.48</b> ±0.26	11000	57.68±0.46	<b>58.90</b> ±0.09
8000	85.51±0.51	<b>85.88</b> ±0.21	12000	58.94±0.80	<b>60.15</b> ±0.33
9000	<b>86.59</b> ±0.34	86.42±0.31	13000	60.06±0.48	<b>60.89</b> ±0.55
10000	<b>87.31</b> ±0.46	87.20±0.29	14000	61.22±0.48	<b>61.70</b> ±0.56

(c) Caltech101			(d) Caltech256		
#Samples	ResNet18 [8]	FR-ResNet18	#Samples	ResNet18 [8]	FR-ResNet18
1000	<b>84.33</b> ±0.70	83.46±0.53	5000	74.75±0.32	<b>75.90</b> ±0.20
2000	90.87±0.38	<b>91.34</b> ±0.17	6000	76.34±0.13	<b>77.30</b> ±0.15
3000	92.78±0.20	<b>93.56</b> ±0.35	7000	77.37±0.23	<b>78.42</b> ±0.26
4000	94.31±0.25	<b>94.49</b> ±0.17	8000	78.06±0.24	<b>79.39</b> ±0.25
5000	94.94±0.28	<b>95.21</b> ±0.20	9000	78.98±0.16	<b>79.69</b> ±0.22
6000	95.42±0.38	<b>95.86</b> ±0.21	10000	79.65±0.21	<b>80.24</b> ±0.09

Table 4: **Comparisons with MLPMixer [2] and ViT [3].** We compare our method to different modern architectures on the CIFAR10 and CIFAR100 datasets in supervised learning.

(a) CIFAR10				
#Samples	ResNet18 [8]	FR-ResNet18	MLPMixer [2]	ViT-B16 [3]
1000	46.18±2.81	<b>53.70</b> ±1.18	49.30±0.57	47.92±0.36
2000	60.36±1.64	<b>70.46</b> ±0.59	59.82±5.40	53.95±0.24
3000	69.22±3.21	<b>76.81</b> ±0.50	61.39±0.51	57.05±0.17
4000	75.45±2.29	<b>79.66</b> ±0.21	64.48±0.93	60.54±0.52
5000	80.23±0.85	<b>82.27</b> ±0.33	66.97±0.57	62.60±0.39
6000	82.52±0.59	<b>83.65</b> ±0.34	68.40±0.71	64.68±0.39
7000	84.08±0.38	<b>84.48</b> ±0.26	69.91±0.67	66.17±0.26
8000	85.51±0.51	<b>85.88</b> ±0.21	71.48±0.37	67.41±0.34
9000	<b>86.59</b> ±0.34	86.42±0.31	73.26±0.55	69.01±0.46
10000	<b>87.31</b> ±0.46	87.20±0.29	73.98±0.62	69.29±0.53

(b) CIFAR100				
#Samples	ResNet18 [8]	FR-ResNet18	MLPMixer [2]	ViT-B16 [3]
5000	38.07±0.92	<b>43.73</b> ±1.36	29.95±0.23	29.28±0.44
6000	41.72±0.61	<b>48.52</b> ±0.78	32.08±0.33	31.60±0.37
7000	46.29±2.07	<b>51.44</b> ±0.45	34.49±0.24	33.03±0.21
8000	50.06±0.60	<b>54.51</b> ±0.40	37.03±0.36	35.06±0.61
9000	53.25±0.32	<b>56.06</b> ±0.40	38.17±0.10	37.08±0.31
10000	55.84±0.90	<b>57.50</b> ±0.56	39.62±0.16	37.57±0.49
11000	57.68±0.46	<b>58.90</b> ±0.09	40.58±0.19	39.87±0.29
12000	58.94±0.80	<b>60.15</b> ±0.33	41.64±0.19	41.68±0.67
13000	60.06±0.48	<b>60.89</b> ±0.55	42.71±0.18	42.36±0.45
14000	61.22±0.48	<b>61.70</b> ±0.56	43.56±0.15	43.62±0.24

Table 5: **Comparisons with Knowledge Distillation baselines.** We compare to several KD methods. Our method outperforms the baselines in the earlier iterations with a large margin.

(a) CIFAR10						
#Samples	ResNet18	FR-ResNet18	DML	KDCL	KD-Student	KD-Teacher
1000	46.18±2.81	<b>53.70±1.18</b>	39.46±4.60	44.19±1.97	46.91±1.09	42.15±1.20
2000	60.36±1.64	<b>70.46±0.59</b>	61.96±1.20	58.87±4.45	59.53±3.19	57.32±6.20
3000	69.22±3.21	<b>76.81±0.50</b>	71.35±1.08	69.68±1.31	67.18±3.75	66.94±1.42
4000	75.45±2.29	<b>79.66±0.21</b>	76.83±1.43	75.98±1.81	75.31±0.64	73.28±3.07
5000	80.23±0.85	<b>82.27±0.33</b>	80.92±0.92	81.14±0.65	78.16±2.87	77.89±1.42
6000	82.52±0.59	83.65±0.34	<b>83.67±0.34</b>	83.57±0.25	82.17±0.88	81.22±1.61
7000	84.08±0.38	84.48±0.26	<b>84.98±0.26</b>	84.66±0.72	83.79±0.37	83.80±0.23
8000	85.51±0.51	85.88±0.21	<b>86.28±0.34</b>	86.17±0.39	85.46±0.37	84.81±0.47
9000	86.59±0.34	86.42±0.31	86.74±0.53	<b>86.98±0.28</b>	86.20±0.28	85.67±0.73
10000	87.31±0.46	87.20±0.29	87.52±0.37	<b>87.55±0.14</b>	87.42±0.27	86.14±0.78
(b) CIFAR100						
#Samples	ResNet18	FR-ResNet18	DML	KDCL	KD-Student	KD-Teacher
5000	38.07±0.92	<b>43.73±1.36</b>	38.79±1.53	39.60±0.91	36.45±2.18	35.77±2.79
6000	41.72±0.61	<b>48.52±0.78</b>	42.17±1.19	44.91±1.16	40.39±0.55	39.74±2.31
7000	46.29±2.07	<b>51.44±0.45</b>	47.28±0.90	48.42±1.18	44.56±1.60	44.37±2.08
8000	50.06±0.60	<b>54.51±0.40</b>	51.88±1.20	52.63±0.31	48.66±1.55	48.11±2.57
9000	53.25±0.32	56.06±0.40	55.13±1.27	<b>56.25±1.09</b>	52.35±1.38	50.94±2.09
10000	55.84±0.90	57.50±0.56	<b>58.08±0.37</b>	57.69±0.72	56.46±1.23	54.32±1.39
11000	57.68±0.46	58.90±0.09	58.88±0.81	<b>60.30±0.81</b>	57.92±0.73	54.95±1.65
12000	58.94±0.80	60.15±0.33	60.30±1.06	<b>61.81±1.00</b>	59.82±0.75	58.38±1.02
13000	60.06±0.48	60.89±0.55	62.58±0.55	<b>63.39±0.73</b>	61.19±0.45	59.59±1.29
14000	61.22±0.48	61.70±0.56	63.55±0.55	<b>64.56±0.56</b>	62.36±1.07	60.80±1.58

Table 6: **Comparison with SimSiam [4].** We compare our method against the CIFAR10 version of SimSiam [4]. Our method significantly outperforms SimSiam in the low-data regime.

#Samples	ResNet18	FR-ResNet18	SimSiam-kNN	SimSiam-Linear
1000	46.18±2.81	<b>53.70±1.18</b>	25.09±1.10	36.67±0.46
2000	60.36±1.64	<b>70.46±0.59</b>	36.53±0.74	45.01±0.11
3000	69.22±3.21	<b>76.81±0.50</b>	48.12±0.33	54.58±0.08
4000	75.45±2.29	<b>79.66±0.21</b>	55.69±0.52	61.39±0.50
5000	80.23±0.85	<b>82.27±0.33</b>	60.96±0.16	65.55±0.40
6000	82.52±0.59	<b>83.65±0.34</b>	65.83±0.12	68.22±0.57
7000	84.08±0.38	<b>84.48±0.26</b>	69.45±0.15	70.53±0.35
8000	85.51±0.51	<b>85.88±0.21</b>	72.34±0.26	73.75±0.11
9000	<b>86.59±0.34</b>	86.42±0.31	74.99±0.25	76.57±0.17
10000	<b>87.31±0.46</b>	87.20±0.29	76.85±0.56	79.20±0.15

Table 7: **Active Learning.** We evaluate the performance of our method on a typical low-data setting, on active learning. We compare to standard maximum entropy-based acquisition and also to two state-of-the-art approaches, to core-set and LLAL [1]. Our method clearly outperforms all the other approaches, especially in the earlier stages.

(a) CIFAR10				
#Samples	ResNet18-Entropy	FR-ResNet18-Entropy	LLAL [1]	ResNet18-core-set
1000	46.28±2.54	<b>53.45±1.30</b>	46.53±1.57	46.52±2.64
2000	60.26±2.95	<b>69.33±2.13</b>	60.96±3.10	62.04±2.57
3000	69.76±2.69	<b>77.64±1.00</b>	73.03±2.68	69.06±2.07
4000	77.62±2.25	<b>81.86±0.62</b>	80.83±0.83	77.33±2.25
5000	81.27±3.14	84.10±0.31	<b>84.14±0.56</b>	80.04±1.80
6000	85.13±0.98	<b>86.33±0.40</b>	85.95±0.50	84.04±0.84
7000	87.05±0.68	<b>87.77±0.36</b>	87.40±0.49	86.24±0.57
8000	88.33±0.64	<b>88.91±0.29</b>	88.55±0.59	87.56±0.49
9000	<b>89.83±0.60</b>	89.59±0.23	89.18±0.41	88.67±0.55
10000	<b>90.40±0.33</b>	90.36±0.23	89.78±0.42	89.86±0.45

(b) CIFAR100				
#Samples	ResNet18-Entropy	FR-ResNet18-Entropy	LLAL [1]	ResNet18-core-set
5000	38.73±0.82	<b>44.22±1.09</b>	37.36±0.92	38.60±0.95
6000	41.95±2.20	<b>48.91±0.66</b>	43.73±0.76	43.62±0.45
7000	45.99±0.74	<b>51.21±0.79</b>	48.26±0.62	48.10±0.86
8000	50.05±0.69	<b>53.80±0.58</b>	51.62±0.60	50.26±1.61
9000	53.70±0.47	<b>56.02±0.63</b>	54.20±0.43	54.57±1.44
10000	55.96±0.54	<b>57.61±0.61</b>	56.18±0.83	56.35±1.41
11000	58.00±1.19	<b>59.73±0.53</b>	57.91±0.57	58.80±0.44
12000	59.90±1.08	<b>60.64±0.49</b>	59.34±0.74	60.43±0.86
13000	61.39±1.37	<b>62.01±0.40</b>	60.78±0.93	61.67±0.79
14000	<b>63.78±0.49</b>	63.27±0.47	61.93±0.71	63.61±0.43

(c) Caltech101				
#Samples	ResNet18-Entropy	FR-ResNet18-Entropy	LLAL [1]	ResNet18-core-set
1000	<b>84.30±0.69</b>	83.46±0.53	83.89±0.33	<b>84.30±0.69</b>
2000	93.27±0.37	93.30±0.33	91.88±0.50	<b>93.50±0.30</b>
3000	95.24±0.27	<b>95.38±0.34</b>	94.02±0.40	95.04±0.17
4000	95.50±0.15	<b>95.83±0.23</b>	93.89±0.23	95.34±0.29
5000	95.42±0.34	<b>95.70±0.38</b>	94.08±0.22	95.56±0.11
6000	95.54±0.17	<b>95.90±0.13</b>	94.18±0.55	95.48±0.14

(d) Caltech256				
#Samples	ResNet18-Entropy	FR-ResNet18-Entropy	LLAL [1]	ResNet18-core-set
5000	74.79±0.33	<b>75.90±0.20</b>	74.06±0.45	74.79±0.33
6000	76.03±0.31	<b>77.11±0.07</b>	74.47±0.21	76.52±0.36
7000	77.53±0.08	<b>78.29±0.19</b>	75.21±0.33	77.44±0.17
8000	78.89±0.28	<b>79.53±0.29</b>	75.92±0.36	78.12±0.22
9000	79.81±0.27	<b>80.34±0.42</b>	76.56±0.19	79.20±0.26
10000	80.50±0.26	<b>81.34±0.34</b>	76.82±0.36	79.93±0.16
11000	81.37±0.39	<b>81.60±0.24</b>	77.12±0.21	80.65±0.27
12000	81.82±0.21	<b>82.44±0.18</b>	77.59±0.25	80.95±0.19
13000	82.22±0.27	<b>82.63±0.15</b>	77.66±0.52	81.62±0.18
14000	82.61±0.31	<b>83.07±0.22</b>	78.14±0.27	81.97±0.25

Table 8: **Backbone Agnosticism.** We evaluate our approach with ResNet34, DenseNet121 [9] and EfficientNet-B3 [10] backbones. Our method is agnostic to the backbone showing benefit in every case.

(a) ResNet34-CIFAR10

#Samples	ResNet34	FR-ResNet34
1000	44.67±1.69	<b>48.88±1.41</b>
2000	60.29±2.04	<b>68.61±1.48</b>
3000	68.68±2.88	<b>76.28±0.69</b>
4000	76.09±1.74	<b>79.73±0.45</b>
5000	79.89±1.11	<b>82.26±0.39</b>
6000	82.80±0.80	<b>83.83±0.20</b>
7000	83.67±0.90	<b>84.99±0.18</b>
8000	85.80±0.32	<b>86.13±0.32</b>
9000	86.53±0.40	<b>86.76±0.15</b>
10000	<b>87.43±0.41</b>	87.33±0.28

(b) ResNet34-CIFAR100

#Samples	ResNet34	FR-ResNet34
5000	32.82±1.83	<b>43.57±1.29</b>
6000	36.14±1.81	<b>48.69±1.04</b>
7000	43.16±1.26	<b>52.30±0.75</b>
8000	47.10±1.91	<b>54.29±0.83</b>
9000	49.95±2.75	<b>55.92±0.39</b>
10000	53.20±1.08	<b>58.11±0.34</b>
11000	55.25±1.86	<b>59.58±0.29</b>
12000	57.19±1.72	<b>60.92±0.41</b>
13000	59.47±1.16	<b>61.64±0.46</b>
14000	60.21±1.29	<b>63.07±0.25</b>

(c) EfficientNetB3-CIFAR10 [10]

#Samples	EfficientNetB3 [10]	FR-EfficientNetB3
1000	25.47±8.65	<b>47.31±1.79</b>
2000	34.42±8.94	<b>59.55±3.52</b>
3000	50.99±3.32	<b>66.67±1.98</b>
4000	59.63±4.99	<b>68.81±2.89</b>
5000	63.51±3.26	<b>72.00±1.27</b>
6000	66.64±5.22	<b>72.96±1.77</b>
7000	67.22±2.53	<b>74.22±2.80</b>
8000	68.36±5.95	<b>74.76±3.14</b>
9000	69.89±4.85	<b>76.61±0.75</b>
10000	69.79±11.33	<b>76.24±1.45</b>

(d) EfficientNetB3-CIFAR100 [10]

#Samples	EfficientNetB3 [10]	FR-EfficientNetB3
5000	13.81±5.56	<b>22.45±1.91</b>
6000	17.59±3.28	<b>23.46±1.50</b>
7000	20.40±2.16	<b>25.22±1.82</b>
8000	22.34±2.23	<b>27.57±2.01</b>
9000	22.24±2.59	<b>27.61±2.93</b>
10000	26.16±2.80	<b>26.89±1.18</b>
11000	24.89±3.60	<b>27.92±1.93</b>
12000	24.61±4.61	<b>30.50±2.07</b>
13000	25.93±2.07	<b>30.87±2.37</b>
14000	26.94±2.85	<b>33.41±2.63</b>

(e) DenseNet121-CIFAR10 [9]

#Samples	DenseNet121 [9]	FR-DenseNet121
1000	42.03±1.44	<b>46.25±0.77</b>
2000	53.43±1.15	<b>58.51±0.74</b>
3000	59.58±1.61	<b>65.85±0.57</b>
4000	64.85±2.26	<b>70.02±0.50</b>
5000	69.23±2.42	<b>73.03±0.54</b>
6000	73.35±0.83	<b>75.18±0.36</b>
7000	75.47±0.60	<b>76.34±0.36</b>
8000	76.96±0.59	<b>77.62±0.42</b>
9000	78.01±0.83	<b>78.56±0.39</b>
10000	78.82±0.65	<b>79.28±0.26</b>

(f) DenseNet121-CIFAR100 [9]

#Samples	DenseNet121 [9]	FR-DenseNet121
5000	27.97±0.72	<b>33.38±0.45</b>
6000	30.61±0.83	<b>36.67±0.72</b>
7000	34.18±0.65	<b>38.99±0.94</b>
8000	35.63±0.65	<b>41.79±0.69</b>
9000	38.45±0.72	<b>43.85±0.81</b>
10000	40.17±0.64	<b>45.54±0.34</b>
11000	42.06±1.01	<b>47.10±0.63</b>
12000	42.92±0.87	<b>48.52±0.23</b>
13000	44.96±0.87	<b>49.09±0.35</b>
14000	46.16±1.70	<b>50.34±0.30</b>

Table 9: **Feature Refiner 9a:** We apply parts of our Feature Refiner step-by-step. First, we use only a single linear layer without any extra activation (512x512 w/o Activation), then apply our dimension reduction step (512x64 w/o Activation). Finally, we evaluate the effect of the LayerNorm layer.

**OJKD ablation 9c:** Our online joint knowledge distillation enables us to utilize the advantages of our Feature Refiner without increasing the number of model parameters at the same time. We also show that we can still improve upon the baseline without our Gradient Gate, but using it gives us extra improvement.

**Number of layers ablation 9b:** We study the effect of the number of nonlinear fully-connected layers. More layers do not lead to better performance.

(a) Feature Refiner

#Samples	ResNet18 [8]	512x512 w/o Activation	512x64 w/o Activation	FR w/o LayerNorm	FR-ResNet18
1000	46.18±2.81	51.09±3.16	52.69±1.01	53.47±1.30	<b>53.70±1.18</b>
2000	60.36±1.64	65.84±3.38	66.16±2.29	69.22±1.04	<b>70.46±0.59</b>
3000	69.22±3.21	73.25±2.65	75.28±0.98	<b>77.12±0.29</b>	76.81±0.50
4000	75.45±2.29	77.39±1.87	79.68±0.31	<b>79.77±0.27</b>	79.66±0.21
5000	80.23±0.85	81.08±0.88	82.03±0.29	82.13±0.25	<b>82.27±0.33</b>
6000	82.52±0.59	83.27±0.90	83.83±0.16	<b>83.94±0.36</b>	83.65±0.34
7000	84.08±0.38	84.59±0.51	84.95±0.38	<b>84.96±0.39</b>	84.48±0.26
8000	85.51±0.51	86.02±0.22	<b>86.27±0.29</b>	86.16±0.38	85.88±0.21
9000	86.59±0.34	86.62±0.38	<b>87.13±0.16</b>	86.69±0.25	86.42±0.31
10000	87.31±0.46	87.65±0.30	<b>87.83±0.28</b>	87.65±0.42	87.20±0.29

(b) Number of layers

#Samples	FR	FR-2layer	FR-3layer	FR-4layer	FR-5layer
1000	<b>53.70±1.18</b>	51.69±1.64	48.11±2.12	47.90±1.64	41.69±3.31
2000	<b>70.46±0.59</b>	68.81±1.42	68.18±2.48	67.83±2.24	64.37±1.73
3000	<b>76.81±0.50</b>	76.31±0.57	75.69±1.77	75.01±2.68	73.61±1.85
4000	<b>79.66±0.21</b>	79.49±0.43	79.29±0.35	78.79±0.42	77.33±1.12
5000	<b>82.27±0.33</b>	81.81±0.42	81.73±0.31	81.36±0.52	80.55±0.50
6000	<b>83.65±0.34</b>	83.49±0.37	83.30±0.48	83.22±0.33	81.89±1.01
7000	<b>84.48±0.26</b>	84.39±0.27	84.28±0.35	83.97±0.37	83.22±0.42
8000	<b>85.88±0.21</b>	85.49±0.22	85.11±0.41	85.21±0.29	84.25±0.27
9000	<b>86.42±0.31</b>	86.06±0.17	85.91±0.30	85.85±0.25	85.18±0.13
10000	<b>87.20±0.29</b>	86.94±0.27	86.64±0.40	86.42±0.29	86.09±0.25

(c) OJKD

#Samples	FR-ResNet18 ExtraHead	FR-ResNet18 OriginalHead	FR-ResNet18 OriginalHead w/o GradGate
1000	53.57±1.02	<b>53.70±1.18</b>	48.44±1.62
2000	70.26±0.61	<b>70.46±0.59</b>	64.89±2.00
3000	76.66±0.55	<b>76.81±0.50</b>	71.45±2.40
4000	79.53±0.24	<b>79.66±0.21</b>	77.35±1.26
5000	82.12±0.35	<b>82.27±0.33</b>	80.84±0.88
6000	83.63±0.29	<b>83.65±0.34</b>	83.06±0.55
7000	84.41±0.26	<b>84.48±0.26</b>	83.97±0.43
8000	85.86±0.26	<b>85.88±0.21</b>	85.67±0.31
9000	86.48±0.32	86.42±0.31	<b>86.59±0.38</b>
10000	<b>87.24±0.26</b>	87.20±0.29	87.18±0.38

Table 10: **Robustness benchmark - Overall.** We evaluate the robustness of our method on the CIFAR10-C dataset. We present the average score over all datasets. Our method performs better in the earlier iterations.

#Samples	ResNet18	FR-ResNet18
1000	40.59±1.65	<b>46.66±1.60</b>
2000	51.46±2.91	<b>58.07±1.99</b>
3000	58.11±2.79	<b>61.45±1.88</b>
4000	61.10±2.27	<b>61.83±1.31</b>
5000	<b>62.80±1.61</b>	61.68±1.75
6000	<b>63.78±1.44</b>	62.82±1.34
7000	<b>64.31±1.36</b>	63.52±1.37
8000	<b>65.50±1.22</b>	63.78±1.72
9000	<b>65.65±1.85</b>	63.38±1.47
10000	<b>65.96±1.19</b>	65.28±1.71

Table 11: **Robustness benchmark - Severity.** We evaluate the robustness of our method on the CIFAR10-C dataset. We present the average score for all severity levels. Our method consistently outperforms the baseline in cases of lower corruption severity. In cases of higher corruption severity (4d and 4e), in the very low-data regime, our method significantly outperforms the baseline. However, with more added data, the baseline starts outperforming our method.

(a) Severity 0			(b) Severity 1			(c) Severity 2		
#Samples	ResNet18	FR-ResNet18	#Samples	ResNet18	FR-ResNet18	#Samples	ResNet18	FR-ResNet18
1000	44.64±1.64	<b>51.82±1.48</b>	1000	40.59±1.65	<b>48.75±1.47</b>	1000	40.65±1.48	<b>46.73±1.45</b>
2000	57.66±3.63	<b>66.03±1.59</b>	2000	51.46±2.91	<b>62.08±1.83</b>	2000	51.75±2.75	<b>58.59±2.01</b>
3000	66.95±2.61	<b>71.58±1.01</b>	3000	58.11±2.79	<b>66.62±1.67</b>	3000	58.67±2.49	<b>62.15±1.88</b>
4000	71.29±1.85	<b>73.39±0.78</b>	4000	61.10±2.27	<b>67.57±1.15</b>	4000	61.81±2.31	<b>62.61±1.39</b>
5000	74.42±1.15	<b>74.90±0.75</b>	5000	62.80±1.61	<b>68.20±1.27</b>	5000	<b>63.58±1.53</b>	62.39±1.80
6000	76.22±0.78	<b>76.37±0.77</b>	6000	63.78±1.44	<b>69.67±1.16</b>	6000	<b>64.83±1.46</b>	63.70±1.41
7000	<b>77.42±0.68</b>	76.91±0.68	7000	64.31±1.36	<b>70.22±1.05</b>	7000	<b>65.30±1.48</b>	64.39±1.40
8000	<b>78.84±0.65</b>	77.93±0.88	8000	65.50±1.22	<b>70.96±1.44</b>	8000	<b>66.88±1.18</b>	64.78±1.85
9000	<b>79.36±0.79</b>	78.11±0.61	9000	65.65±1.85	<b>70.86±1.15</b>	9000	<b>67.01±1.99</b>	64.34±1.51
10000	<b>79.94±0.71</b>	79.45±0.85	10000	65.96±1.19	<b>72.64±1.33</b>	10000	<b>67.30±1.31</b>	66.45±1.75

(d) Severity 3			(e) Severity 4		
#Samples	ResNet18	FR-ResNet18	#Samples	ResNet18	FR-ResNet18
1000	39.08±1.71	<b>44.72±1.64</b>	1000	36.62±1.94	<b>41.30±1.94</b>
2000	48.93±2.59	<b>54.70±2.17</b>	2000	44.75±2.42	<b>48.94±2.35</b>
3000	54.30±3.21	<b>57.11±2.31</b>	3000	48.16±3.44	<b>49.78±2.55</b>
4000	56.47±2.56	<b>56.89±1.57</b>	4000	<b>49.63±2.62</b>	48.71±1.66
5000	<b>57.64±1.91</b>	55.99±2.32	5000	<b>49.66±2.16</b>	46.93±2.59
6000	<b>58.24±1.79</b>	56.87±1.64	6000	<b>49.36±2.02</b>	47.48±1.71
7000	<b>58.51±1.75</b>	57.86±1.85	7000	<b>49.29±1.83</b>	48.23±1.86
8000	<b>59.54±1.67</b>	57.62±2.22	8000	<b>49.52±1.88</b>	47.61±2.21
9000	<b>59.62±2.48</b>	56.98±1.96	9000	<b>49.23±2.64</b>	46.60±2.12
10000	<b>59.83±1.43</b>	59.23±2.28	10000	<b>49.38±1.47</b>	48.63±2.32

Table 12: **Robustness benchmark - Corruption types.** We evaluate the robustness of our method on the CIFAR10-C dataset. We present the average score for all corruption types. Our method consistently outperforms the baseline in the earlier iterations.

(a) Brightness		(b) Contrast		(c) Defocus Blur		(d) Elastic Transform		
#Samples	ResNet18	FR-ResNet18	ResNet18	FR-ResNet18	ResNet18	FR-ResNet18	ResNet18	FR-ResNet18
1000	37.59±2.84	<b>48.20</b> ±1.32	23.42±1.61	<b>28.70</b> ±1.62	42.95±1.59	<b>47.82</b> ±1.56	42.78±0.80	<b>47.80</b> ±1.18
2000	54.49±5.07	<b>64.82</b> ±1.91	30.62±3.60	<b>41.13</b> ±2.11	54.19±2.76	<b>59.36</b> ±2.37	54.17±2.68	<b>59.62</b> ±1.93
3000	67.14±3.91	<b>73.69</b> ±0.93	40.97±5.18	<b>50.98</b> ±2.13	59.96±2.34	<b>62.97</b> ±1.31	61.24±1.07	<b>63.86</b> ±0.87
4000	72.49±2.30	<b>76.91</b> ±0.78	46.10±3.49	<b>55.30</b> ±1.82	<b>63.63</b> ±1.77	63.44±1.05	65.28±1.33	<b>65.30</b> ±0.73
5000	76.63±1.46	<b>79.56</b> ±0.41	50.80±2.87	<b>56.17</b> ±2.15	<b>65.42</b> ±1.09	64.29±1.92	<b>68.09</b> ±0.92	67.12±1.18
6000	79.65±0.50	<b>81.19</b> ±0.35	55.69±1.01	<b>59.32</b> ±1.24	<b>66.52</b> ±1.42	64.81±1.22	<b>69.94</b> ±0.87	68.30±0.82
7000	81.26±0.44	<b>82.17</b> ±0.40	58.19±1.81	<b>60.35</b> ±0.99	<b>68.72</b> ±1.00	67.48±1.52	<b>71.78</b> ±0.77	69.71±0.94
8000	83.11±0.56	<b>83.50</b> ±0.30	61.30±1.16	<b>61.58</b> ±1.35	<b>69.99</b> ±1.03	69.85±1.16	<b>73.03</b> ±0.70	70.92±0.81
9000	84.00±0.74	<b>84.18</b> ±0.41	61.99±2.51	<b>62.83</b> ±1.04	<b>71.22</b> ±1.86	70.25±1.20	<b>74.45</b> ±0.92	71.94±0.52
10000	<b>85.03</b> ±0.54	85.01±0.41	<b>63.70</b> ±1.43	62.84±1.17	<b>71.81</b> ±0.92	70.62±1.74	<b>74.83</b> ±0.55	72.18±1.25
(e) Fog		(f) Frost		(g) Gaussian Blur		(h) Gaussian Noise		
#Samples	ResNet18	FR-ResNet18	ResNet18	FR-ResNet18	ResNet18	FR-ResNet18	ResNet18	FR-ResNet18
1000	31.14±1.49	<b>37.96</b> ±2.24	31.10±2.75	<b>42.01</b> ±1.40	41.41±1.94	<b>45.19</b> ±1.72	44.93±1.33	<b>50.12</b> ±2.26
2000	42.19±5.44	<b>56.09</b> ±2.64	46.39±4.69	<b>56.12</b> ±2.13	51.61±2.40	<b>55.08</b> ±2.56	53.56±1.83	<b>58.23</b> ±1.91
3000	56.09±5.01	<b>66.27</b> ±1.82	56.93±1.86	<b>62.03</b> ±1.00	55.34±2.73	<b>57.22</b> ±1.55	55.22±3.72	<b>55.53</b> ±3.69
4000	62.36±3.69	<b>70.24</b> ±1.44	60.45±2.12	<b>62.64</b> ±1.19	<b>58.39</b> ±1.97	56.24±1.35	<b>54.29</b> ±3.50	51.91±1.88
5000	67.62±2.53	<b>72.08</b> ±1.22	63.01±1.34	62.82±1.26	<b>59.07</b> ±1.28	60.02±2.63	<b>52.94</b> ±2.21	46.57±3.13
6000	71.34±0.69	<b>74.88</b> ±0.36	<b>64.00</b> ±1.15	63.70±1.16	<b>59.41</b> ±1.89	56.19±1.79	<b>50.95</b> ±2.13	46.62±1.81
7000	73.76±1.00	<b>75.69</b> ±0.48	<b>63.95</b> ±1.46	62.90±0.85	<b>61.62</b> ±1.25	59.34±2.14	<b>47.14</b> ±2.94	45.87±2.64
8000	76.34±0.70	<b>77.01</b> ±1.10	<b>65.62</b> ±0.73	63.62±1.82	61.97±1.58	<b>62.23</b> ±1.51	<b>46.82</b> ±2.47	40.84±3.51
9000	76.93±1.62	<b>77.41</b> ±0.49	<b>65.71</b> ±1.17	63.22±1.36	<b>63.54</b> ±2.43	61.63±1.79	<b>43.21</b> ±3.16	37.68±2.46
10000	<b>78.23</b> ±0.58	77.86±0.80	65.36±1.53	<b>65.45</b> ±1.81	<b>63.62</b> ±1.28	61.97±2.31	42.65±2.39	<b>44.18</b> ±3.70
(i) Glass Blur		(j) Impulse Noise		(k) JPEG Compression		(l) Motion Blur		
#Samples	ResNet18	FR-ResNet18	ResNet18	FR-ResNet18	ResNet18	FR-ResNet18	ResNet18	FR-ResNet18
1000	43.52±0.92	<b>47.97</b> ±1.02	43.73±1.68	<b>48.62</b> ±2.10	45.99±1.64	<b>52.72</b> ±1.68	41.28±0.73	<b>44.41</b> ±1.43
2000	51.15±1.17	<b>53.77</b> ±1.82	50.82±2.62	<b>56.10</b> ±1.58	58.20±3.15	<b>65.35</b> ±1.39	49.58±1.35	<b>52.20</b> ±2.18
3000	<b>52.75</b> ±4.57	47.76±4.50	53.72±2.62	<b>56.79</b> ±2.51	66.52±1.47	<b>69.30</b> ±0.66	52.84±1.76	<b>55.48</b> ±1.40
4000	<b>52.87</b> ±2.81	45.57±1.19	54.97±2.36	<b>57.51</b> ±1.63	<b>70.34</b> ±1.12	69.76±0.75	55.21±1.45	<b>55.32</b> ±0.97
5000	<b>48.92</b> ±3.46	40.31±3.59	<b>54.94</b> ±2.68	54.15±1.83	<b>73.16</b> ±0.40	71.36±0.92	<b>56.10</b> ±1.37	55.93±1.96
6000	<b>45.82</b> ±3.12	40.80±3.23	54.43±2.42	<b>54.49</b> ±2.36	<b>74.11</b> ±0.41	71.42±0.67	57.58±1.83	<b>58.28</b> ±1.55
7000	<b>42.95</b> ±2.47	39.82±1.69	53.54±1.30	<b>56.26</b> ±2.16	<b>74.65</b> ±0.75	71.43±1.14	<b>61.26</b> ±1.11	59.77±1.28
8000	<b>42.23</b> ±1.01	36.06±3.74	<b>54.80</b> ±2.06	50.89±3.21	<b>74.64</b> ±0.91	71.74±0.71	<b>63.72</b> ±0.68	62.99±2.29
9000	<b>41.94</b> ±3.82	35.76±2.72	<b>52.86</b> ±2.68	49.95±3.26	<b>75.44</b> ±0.83	72.02±0.36	<b>64.79</b> ±1.58	63.01±2.38
10000	<b>41.80</b> ±1.89	40.06±1.61	53.27±1.23	<b>54.26</b> ±2.48	<b>75.12</b> ±0.74	72.55±0.82	<b>66.39</b> ±0.85	63.87±2.22
(m) Pixelate		(n) Saturate		(o) Shot Noise		(p) Snow		
#Samples	ResNet18	FR-ResNet18	ResNet18	FR-ResNet18	ResNet18	FR-ResNet18	ResNet18	FR-ResNet18
1000	45.22±1.42	<b>52.40</b> ±1.40	39.36±2.22	<b>47.23</b> ±2.06	45.43±1.46	<b>50.87</b> ±2.12	39.21±2.46	<b>47.02</b> ±0.67
2000	57.97±3.07	<b>63.62</b> ±1.82	52.11±4.35	<b>62.93</b> ±1.72	55.39±2.06	<b>61.13</b> ±2.02	51.42±2.25	<b>57.71</b> ±1.91
3000	<b>64.90</b> ±3.04	64.38±2.37	63.23±3.88	<b>71.16</b> ±1.44	58.90±2.99	<b>60.46</b> ±2.98	58.89±1.22	<b>62.09</b> ±1.77
4000	<b>66.95</b> ±2.42	63.82±1.79	69.32±2.68	<b>74.67</b> ±0.82	<b>59.20</b> ±3.05	57.72±1.80	62.17±1.50	<b>62.64</b> ±1.26
5000	<b>69.54</b> ±1.37	66.11±1.55	73.39±1.15	<b>77.36</b> ±0.74	<b>58.82</b> ±1.55	53.77±2.49	<b>64.21</b> ±1.29	63.95±0.89
6000	<b>69.85</b> ±1.83	66.15±1.20	76.89±0.75	<b>79.39</b> ±0.81	<b>57.89</b> ±1.92	54.55±1.57	65.42±1.18	<b>65.79</b> ±1.05
7000	<b>69.99</b> ±1.02	65.58±1.60	79.20±0.59	<b>80.58</b> ±0.46	<b>55.05</b> ±2.56	54.10±2.38	<b>65.80</b> ±1.23	65.52±0.56
8000	67.64±2.17	<b>67.71</b> ±1.40	81.37±0.62	<b>81.92</b> ±0.44	<b>56.13</b> ±2.39	51.48±3.16	<b>68.31</b> ±0.55	67.19±1.13
9000	<b>69.66</b> ±2.17	67.06±2.22	82.31±0.73	<b>82.83</b> ±0.37	<b>53.03</b> ±2.62	48.05±2.20	<b>69.31</b> ±0.99	67.42±1.05
10000	<b>69.21</b> ±1.25	68.16±1.47	83.09±0.47	<b>83.09</b> ±0.35	52.35±2.08	<b>53.83</b> ±3.34	<b>69.41</b> ±1.23	69.12±0.91
(q) Spatter		(r) Speckle Noise		(s) Zoom Blur				
#Samples	ResNet18	FR-ResNet18	ResNet18	FR-ResNet18	ResNet18	FR-ResNet18	ResNet18	FR-ResNet18
1000	45.47±1.32	<b>51.61</b> ±1.21	45.36±1.52	<b>50.82</b> ±2.03	41.37±1.69	<b>45.15</b> ±1.31		
2000	56.59±2.80	<b>63.81</b> ±1.02	55.40±2.23	<b>60.97</b> ±2.11	51.85±2.12	<b>55.22</b> ±2.67		
3000	64.52±1.91	<b>69.56</b> ±0.66	58.96±2.54	<b>60.78</b> ±2.86	56.05±1.60	<b>57.18</b> ±1.35		
4000	68.67±1.46	<b>71.30</b> ±1.04	<b>59.23</b> ±3.01	58.18±1.95	<b>59.03</b> ±1.56	56.37±1.49		
5000	71.39±1.39	<b>72.80</b> ±0.38	<b>58.98</b> ±1.42	54.54±2.32	<b>60.11</b> ±1.15	56.99±2.60		
6000	73.47±0.84	<b>74.01</b> ±0.88	<b>58.44</b> ±2.01	55.58±1.58	<b>60.39</b> ±1.63	58.10±1.72		
7000	<b>74.33</b> ±0.71	74.24±0.82	<b>55.74</b> ±2.54	55.50±2.30	<b>63.03</b> ±1.18	60.61±1.61		
8000	<b>75.83</b> ±0.45	75.07±0.62	<b>57.57</b> ±2.40	53.61±3.04	<b>64.07</b> ±1.18	63.56±1.41		
9000	<b>75.83</b> ±0.91	75.37±0.75	<b>54.78</b> ±2.46	50.15±2.10	<b>66.30</b> ±2.24	63.39±1.24		
10000	<b>76.75</b> ±0.67	75.43±0.77	54.24±1.95	<b>55.86</b> ±2.98	<b>66.41</b> ±1.19	63.98±2.29		

## References

- [1] Donggeun Yoo and In So Kweon. Learning loss for active learning. In *CVPR*, 2019. [1](#), [2](#), [8](#)
- [2] Ilya O. Tolstikhin, Neil Houlsby, Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Thomas Unterthiner, Jessica Yung, Andreas Steiner, Daniel Keysers, Jakob Uszkoreit, Mario Lucic, and Alexey Dosovitskiy. Mlp-mixer: An all-mlp architecture for vision. In *NeurIPS*, 2021. [2](#), [6](#)
- [3] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021. [2](#), [6](#)
- [4] Tianyu Hua et al. Simsiam. <https://github.com/PatrickHua/SimSiam>, 2021. [2](#), [7](#)
- [5] Mingjian Zhu, Kai Han, Yehui Tang, and Yunhe Wang. Visual transformer pruning. *CoRR*, abs/2104.08500, 2021. URL <https://arxiv.org/abs/2104.08500>. [2](#)
- [6] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015. [3](#)
- [7] Dan Hendrycks and Thomas G. Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019. URL <https://openreview.net/forum?id=HJz6ticCqYm>.
- [8] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. [6](#), [10](#)
- [9] Gao Huang, Zhuang Liu, Laurens van der Maaten, and Kilian Q. Weinberger. Densely connected convolutional networks. In *CVPR*, 2017. [9](#)
- [10] Mingxing Tan and Quoc V. Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *ICML*, 2019. [9](#)