Label Noise Resiliency with Self-supervised Representations

Zahra Vaseqi, Ibtihel Amara, Samrudhdhi Rangrej
{zahra.vaseqi;ibtihel.amara;samrudhdhi.rangrej}@mail.mcgill.ca

Motivation

- Label noise leads to inaccurate modeling and poor generalization.
- Solutions include pretraining the feature extractor using:
  Self-supervised learning (SSL) without labels
  Or
  Supervised learning with clean in-domain dataset, i.e. transfer learning (TL)
- We compare noise resiliency of <u>five</u> SSL methods against transfer learning.

Noise Types

- *Symmetric* - label is corrupted uniformly and randomly using one of the incorrect classes.
- *Asymmetric* - label is corrupted in a class-dependent manner with one of the incorrect classes.

Experiment 1: Comparison of five SSL methods.

- SimCLR, MoCo, BYOL, SimSiam, and SwaV.

<u>Training</u>

- Pretrain each SSL model with ResNet-18 backbone on CIFAR-10 training set without labels.
- Train final linear classifier on CIFAR-10 training set but with noise-corrupted labels
  - Remove the projection head and freeze the backbone
  - Normalize the output of the backbone
  - Add a linear classifier initialized with random weights.
  - Train the classifier on noise-corrupted set

<u>Testing</u>

- Test the classifier on the clean test set and evaluate top-1 accuracy.

<u>Ablation</u>: Test effect of temperature in InfoNCE loss.

Results (Experiment 1)

- MoCo and SimCLR have the highest accuracies, respectively.
- SimCLR is the most robust method
  performance drop of 1.8% for asym noise and 4.8% for sym noise
- MoCo is the least robust method
  performance drop of 3.5% for asym noise and 8% for sym noise

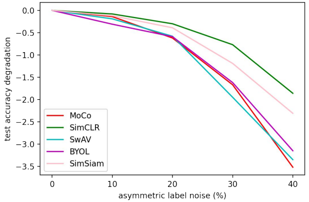

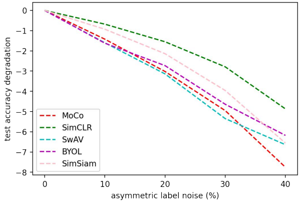

Noise	Symmetric Noise					Asymmetric Noise				
(%)	MoCo	SimCLR	SwAV	SimSiam	BYOL	MoCo	SimCLR	SwAV	SimSiam	BYOL
00	83.57	81.92	66.71	80.46	67.29	83.57	81.92	66.71	80.46	67.29
10	82.15	81.24	65.11	79.53	65.67	83.43	81.84	66.52	80.34	66.98
20	80.55	80.37	63.57	78.32	64.56	82.95	81.62	66.13	80.07	66.70
30	78.62	79.13	61.36	76.50	62.65	81.90	81.15	64.76	79.27	65.67
40	75.82	77.05	60.07	73.90	61.11	80.05	80.06	63.36	78.15	64.14
50	73.26	75.24	57.82	72.28	59.15	76.77	77.96	61.37	74.82	62.29
60	70.24	71.56	55.25	68.49	57.00	73.12	73.91	59.11	72.13	59.20
70	67.28	70.06	54.33	65.12	55.12	69.50	70.24	55.40	67.44	56.28
80	62.65	65.71	51.16	62.42	52.89	62.87	65.05	51.97	62.26	51.74
90	57.62	60.38	47.41	57.45	47.79	56.38	57.40	46.82	55.51	47.09

Table 1: Linear Classification test accuracy(%) on noisy CIFAR-10 using Resnet-18 backbone.

- SimCLR feat distribution is tighter →high tolerance to noise
- MoCo feat distribution is scattered→low tolerance to noise

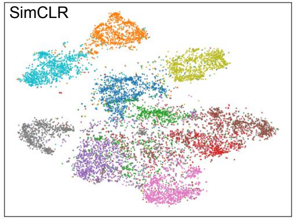

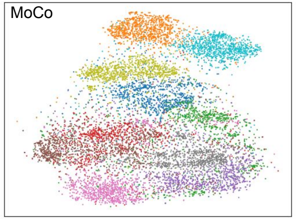

Experiment 2: Comparison of SSL pretraining vs TL.

<u>Data</u>

- We split the CIFAR-10 training set into two equal subsets:
  - Subset 1 (clean labels) - used for pretraining
  - Subset 2 (corrupted labels) - used for finetuning

<u>Training</u>

- Pretrain feature extractor on subset 1 using:
  - SSL (MoCo and SimCLR) without labels.
  - TL with clean labels.
- Train final linear classifier on subset 2.
  - Remove the projection head and freeze the backbone
  - Normalize the output of the backbone
  - Add a linear classifier initialized with random weights
  - Train the classifier on the noise-corrupted set

<u>Testing</u>

- Test the classifier on the clean test set and evaluate top-1 accuracy.

Results (Experiment 2)

- TL outperforms SSL in low noise regime.
- SSL outperforms TL in high noise regime.

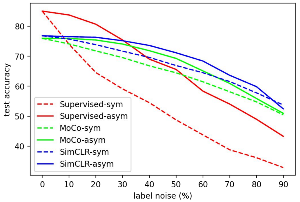

Conclusions

- Robustness to asymmetric noise > symmetric noise.
- SimCLR/MoCo achieves the most/least robustness.
- Tuning temperature in InfoNCE loss improves noise resilience.
- SSL is more noise resilient than TL.