AugLoss: A Robust, Reliable Methodology for Real-World Corruptions

Kyle Otstot[1], John Kevin Cava[1], Tyler Sypherd[1], Lalitha Sankar[1]

[1]Arizona State University, **{kotstot, jcava, tsypherd, lsankar}@asu.edu**

ICML International Conference On Machine Learning

Image Classification Problem

Given random variables $(X, Y) \sim q_{X,Y}$ goal for our model is to learn from the train set a good estimator of $q_{Y|X}$.

This setup is biased to the **i.i.d.** assumptions, *i.e.* the train examples are distributed identically to the examples encountered in the wild.

What if this assumption isn't true? Do our models still learn a good estimate?

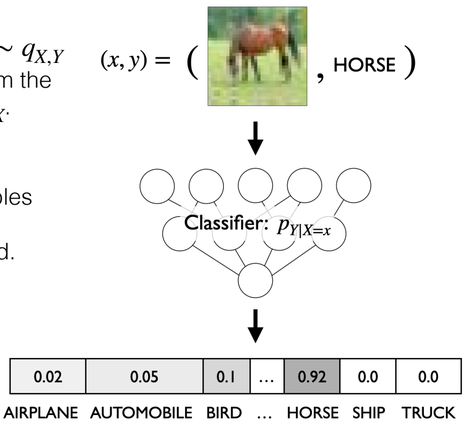

Dataset Corruption

The train set is drawn from a misaligned $\tilde{q}_{X,Y} = \tilde{q}_{Y|X} \cdot \tilde{q}_X$

$\tilde{q}_{Y|X}$: corruption of the true posterior

$\tilde{q}_X$: corruption of the true prior

Train-time noisy labeling

Real world labels are 8-38% noisy **[1]**

Test-time domain shift

Classifiers vulnerable to small image corruptions **[2]**

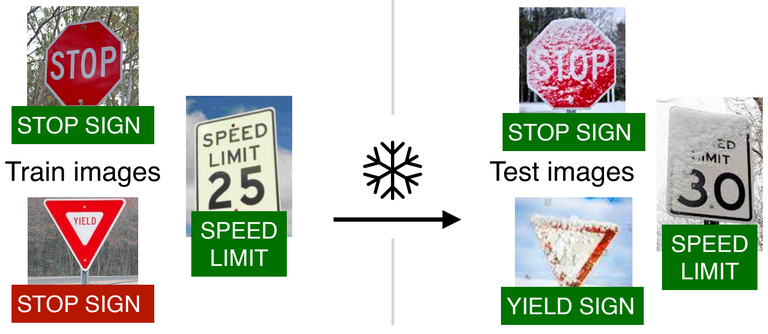

Existing Remedies

Robust Loss Functions for noisy labeling

*e.g., F*ocal loss **[3]**

$$\ell(\hat{p}, y; \gamma) = \left(1 - \hat{p}(y)\right)^{\gamma} \log \hat{p}(y)$$

NCE+RCE **[4]**

$$\ell(\hat{p}, y; \beta_1, \beta_2) = \beta_1 \cdot \frac{\log \hat{p}(y)}{\sum_k \log \hat{p}(k)} + \beta_2 \cdot \left(1 - \hat{p}(y)\right)$$

and α-loss **[5]**

$$\ell(\hat{p}, y; \alpha) = \frac{\alpha}{\alpha - 1}\left(1 - \hat{p}(y)^{1 - \frac{1}{\alpha}}\right)$$

All empirically shown to outperform *cross entropy* loss under label noise.

Data Augmentation for domain adaptation

e.g., AugMix with Jensen-Shannon Divergence consistency loss **[6]**

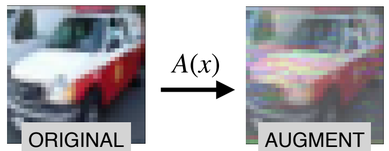

$$D_{JS}(\hat{P}) = \frac{1}{|\hat{P}|} \sum_{\hat{p}_i \in \hat{P}} D_{KL}\left(\hat{p}_i \| \hat{p}_{mix}\right)$$

State-of-the-art results on CIFAR 10/100-C and ImageNet-C

Can we leverage these well-studied techniques to address both types of dataset corruption *simultaneously*?

AugLoss Framework

The *unification* of **data augmentation** and **robust loss functions** to combat noisy labeling at train time and out-of-distribution features at test time.

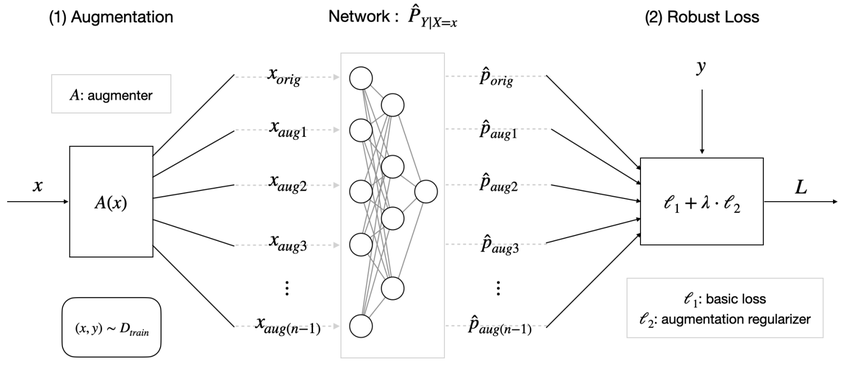

Experiments on CIFAR-10/100

Train label noise generation

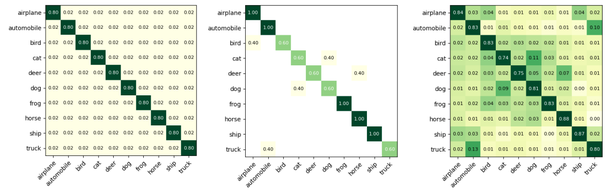

Symmetric Asymmetric CIFAR-10N **[7]**

Test images: CIFAR-10/100-C [1]

CIFAR-10N Results:

Augment	Loss	CIFAR-10N Aggregate	Random 1	Random 2	Random 3	Worst
NoAug	CE	32.24 ± 0.41	37.56 ± 0.18	37.66 ± 0.30	37.96 ± 0.13	49.25 ± 0.34
	Focal	29.85 ± 0.42	34.84 ± 0.46	34.85 ± 0.52	35.20 ± 0.39	48.05 ± 0.96
	NCE+RCE	30.18 ± 0.21	31.11 ± 0.73	31.49 ± 0.31	32.35 ± 1.80	$\mathbf{38.13 \pm 0.46}$
	α-loss	$\mathbf{29.22 \pm 0.79}$	$\mathbf{30.71 \pm 1.18}$	$\mathbf{30.44 \pm 0.88}$	$\mathbf{31.34 \pm 0.36}$	39.93 ± 0.35
AugMix	CE	15.40 ± 0.30	18.59 ± 0.15	18.76 ± 0.19	18.95 ± 0.17	29.73 ± 0.28
	Focal	13.28 ± 0.16	$\mathbf{13.77 \pm 0.11}$	$\mathbf{13.60 \pm 0.30}$	$\mathbf{13.61 \pm 0.20}$	24.31 ± 0.18
	NCE+RCE	13.72 ± 0.27	14.16 ± 0.03	13.85 ± 0.18	14.07 ± 0.09	$\mathbf{18.14 \pm 0.32}$
	α-loss	$\mathbf{13.06 \pm 0.13}$	14.07 ± 0.28	14.04 ± 0.07	14.00 ± 0.06	21.25 ± 0.04

***AugLoss* best addresses dataset corruption**

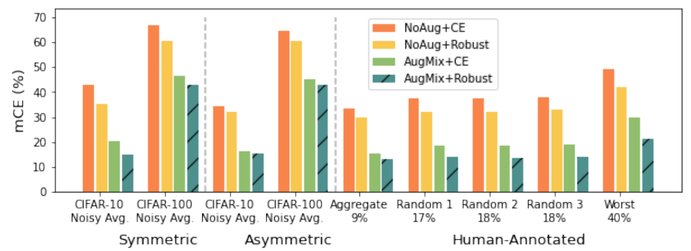

No single loss function works best across all settings

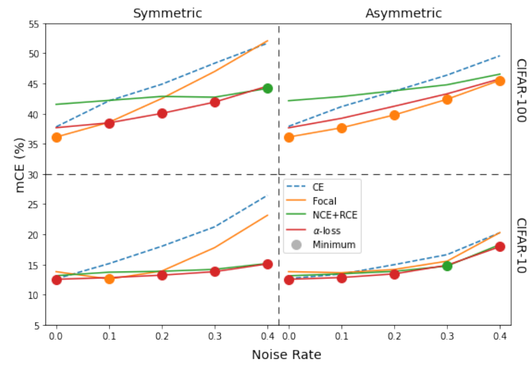

References

[1] H. Song, M. Kim, and J.-G. Lee, "SELFIE: Refurbishing unclean samples for robust deep learning," in *Proceedings of the 36th International Conference on Machine Learning*.
[2] D. Hendrycks and T. Dietterich, "Benchmarking neural network robustness to common corruptions and perturbations," 2019.
[3] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollar, "Focal loss for dense object detection," *2017 IEEE International Conference on Computer Vision (ICCV)*, Oct 2017.
[4] X. Ma, H. Huang, Y. Wang, S. Romano, S. Erfani, and J. Bailey, "Normalized loss functions for deep learning with noisy labels," 2020.
[5] T. Sypherd, M. Diaz, J. K. Cava, G. Dasarathy, P. Kairouz, and L. Sankar, "A tunable loss function for robust classification: Calibration, landscape, and generalization," 2021.
[6] D. Hendrycks, N. Mu, E. D. Cubuk, B. Zoph, J. Gilmer, and B. Lakshminarayanan, "Augmix: A simple data processing method to improve robustness and uncertainty," 2019.
[7] J. Wei, Z. Zhu, H. Cheng, T. Liu, G. Niu, and Y. Liu, "Learning with noisy labels revisited: A study using real-world human annotations," 2021.

arXiv GitHub