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

K. Otstot[1] J.K. Cava[1] T. Sypherd[1] L. Sankar[1]

[1] School of Electrical, Computer, and Energy Engineering
Arizona State University

Arizona State University

ICML PODS Workshop, July 2022

International Conference On Machine Learning

Introduction

- **Image Classification**
  - Given feature-label pair of random variables $(X, Y) \sim q_{X,Y}$, goal for the model is to learn a classifier that approximates $q_{Y|X}$
  - Model learns from dataset drawn from $q_{X,Y}$, the underlying joint distribution – I.I.D. assumption **[1]**
  - Problem: what if the dataset is "**corrupted**", i.e. drawn from a misaligned joint distribution $\tilde{q}_{X,Y}$?

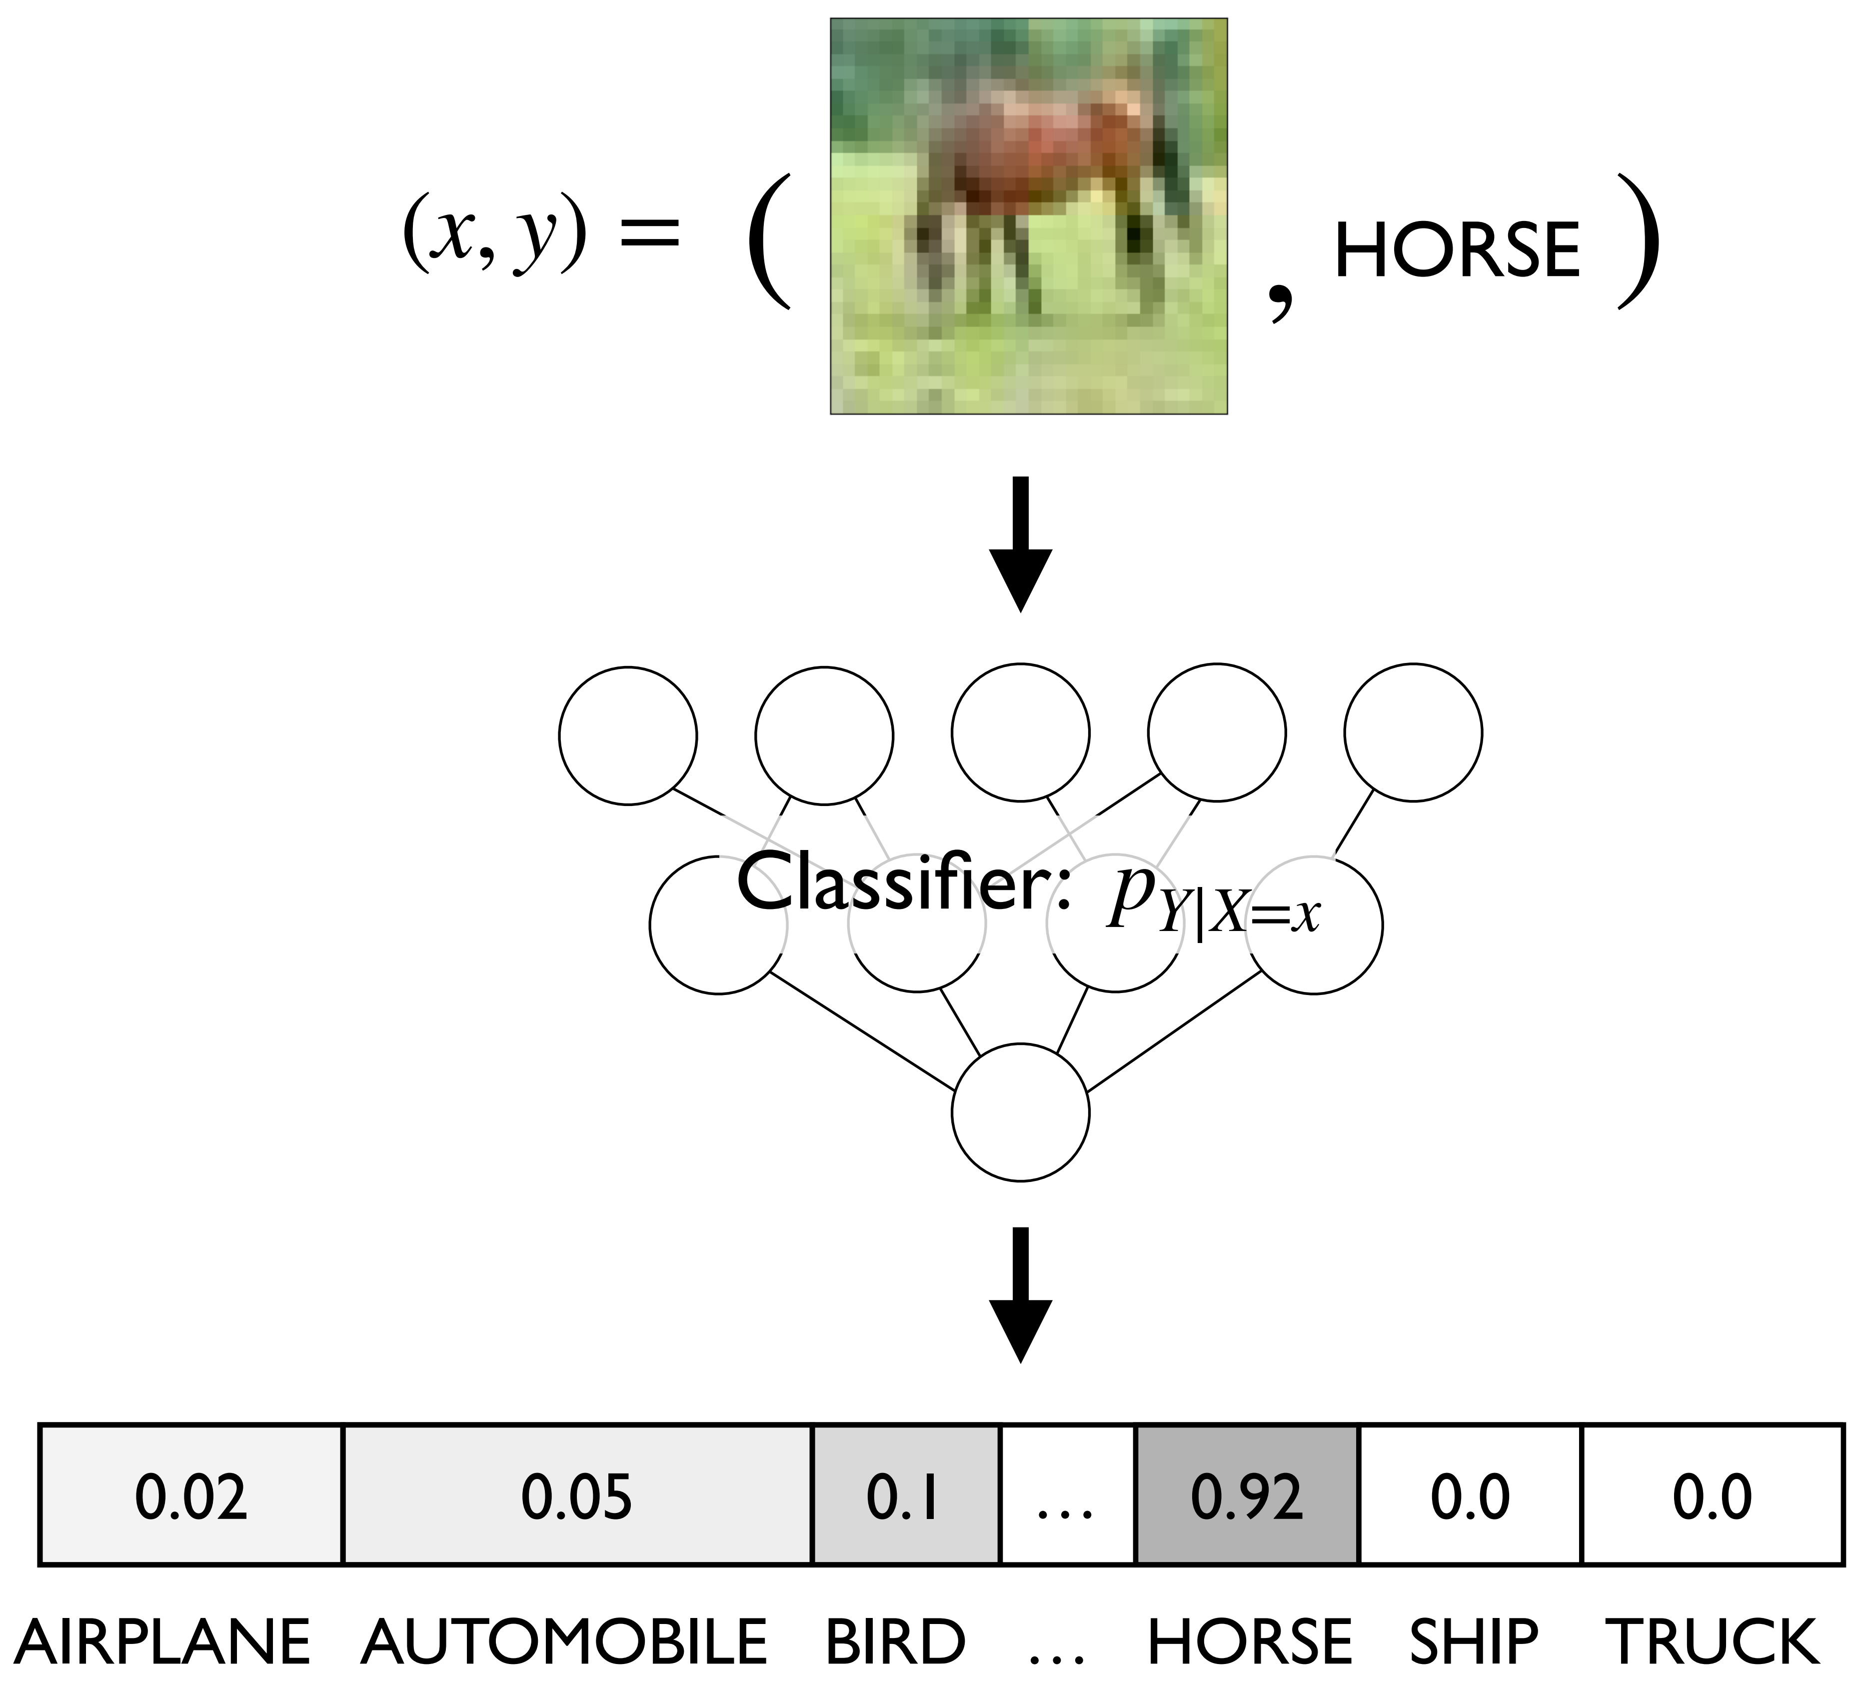

Introduction

- **Dataset Corruption**
  - Dataset is drawn from $\tilde{q}_{X,Y} = \tilde{q}_{Y|X} \cdot \tilde{q}_X$
  - $\tilde{q}_{Y|X}$: **corruption of the true posterior**
    - Approximately 8-38% of labels in real-world datasets are noisy [2]
    - Flaws in data collection, e.g. crowdsourcing [3]
  - $\tilde{q}_X$: **corruption of the true prior**
    - Test-time feature distribution shifts
    - Small corruptions to test images can subvert existing classifiers [4]

Train images:

Test images:

Related Work

- **Robust Loss Functions**
  - A proposed remedy for **noisy labeling** in the train data
  - Cross entropy (CE) loss shown to be non-robust under label noise [5]
  - Focal loss [6], NCE+RCE loss [5], and α-loss [7] have all been experimentally shown to outperform CE loss under label noise
- **Data Augmentation**
  - A proposed remedy for **test-time feature distribution shifts**
  - AugMix [8] has achieved state-of-the-art results on CIFAR-10/100-C

AugLoss Framework

- ***AugLoss:*** our learning methodology unifying data augmentation and robust loss functions to combat both noisy labeling and distribution shifts

<u>Important settings</u>

1. Augmentation technique (augmenter + regularizer)

2. Neural network model

3. Robust (basic) loss function

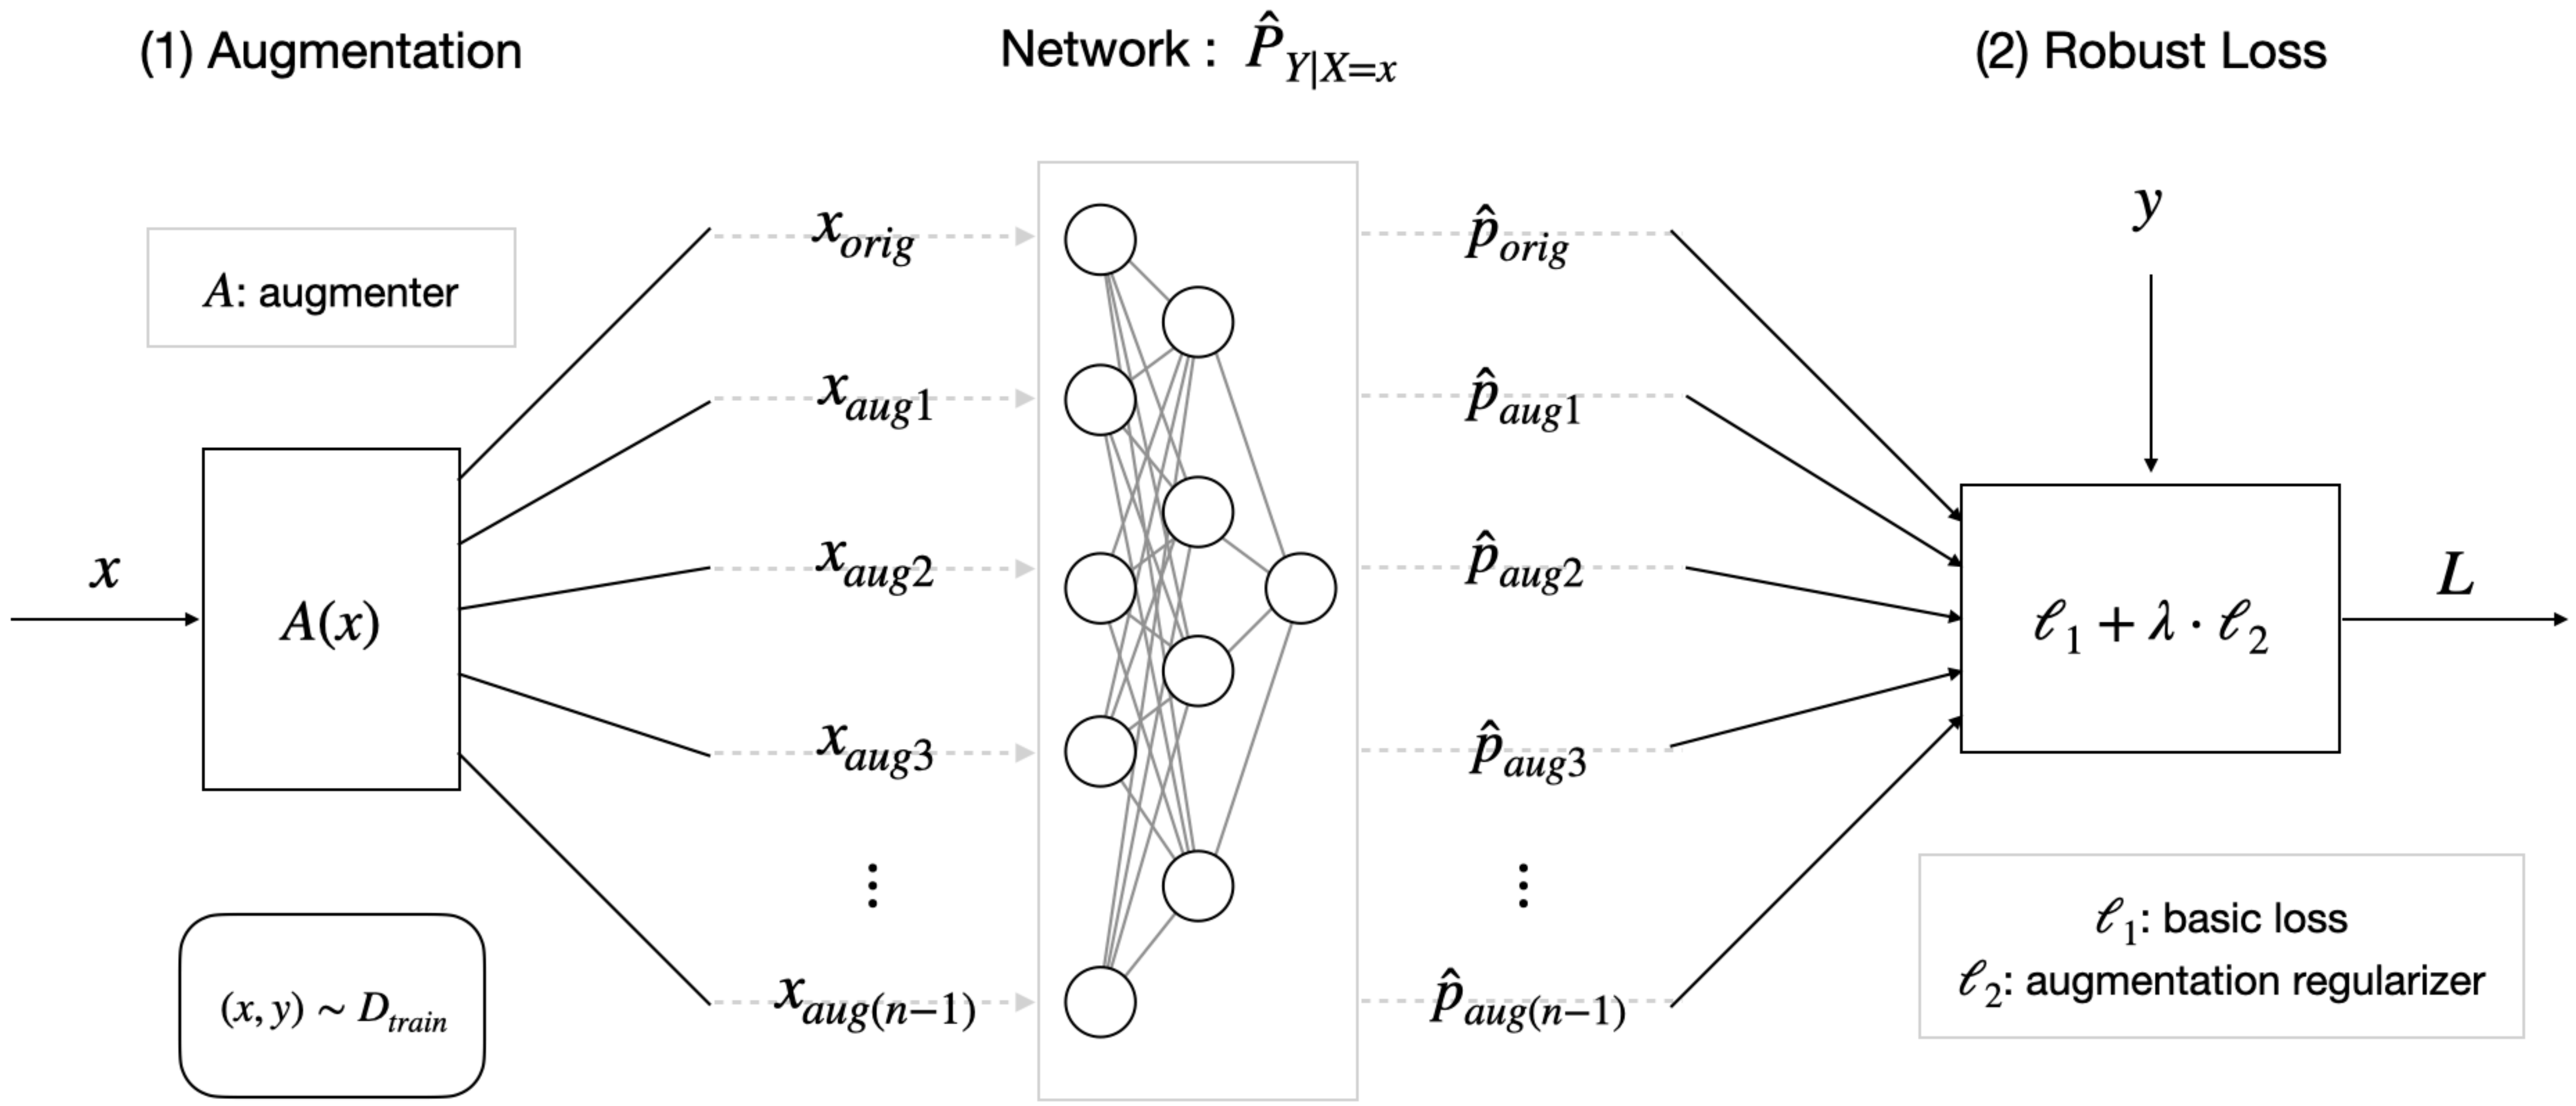

Experiments

- **Question:** How do *AugLoss*-specific methods perform under settings of noisy labeling and distribution shifts, compared to previous state-of-the-art approaches?

- **Datasets:** CIFAR-10 and CIFAR-100
- **Label noise generation:** synthetic (symmetric, asymmetric) and human-annotated (CIFAR-N [9])
- **Distribution shift modeling:** train on traditional (clean) CIFAR, evaluate on CIFAR-C [4]
  - **Performance metric:** mean corruption error (mCE) across the 15 corruptions in CIFAR-C

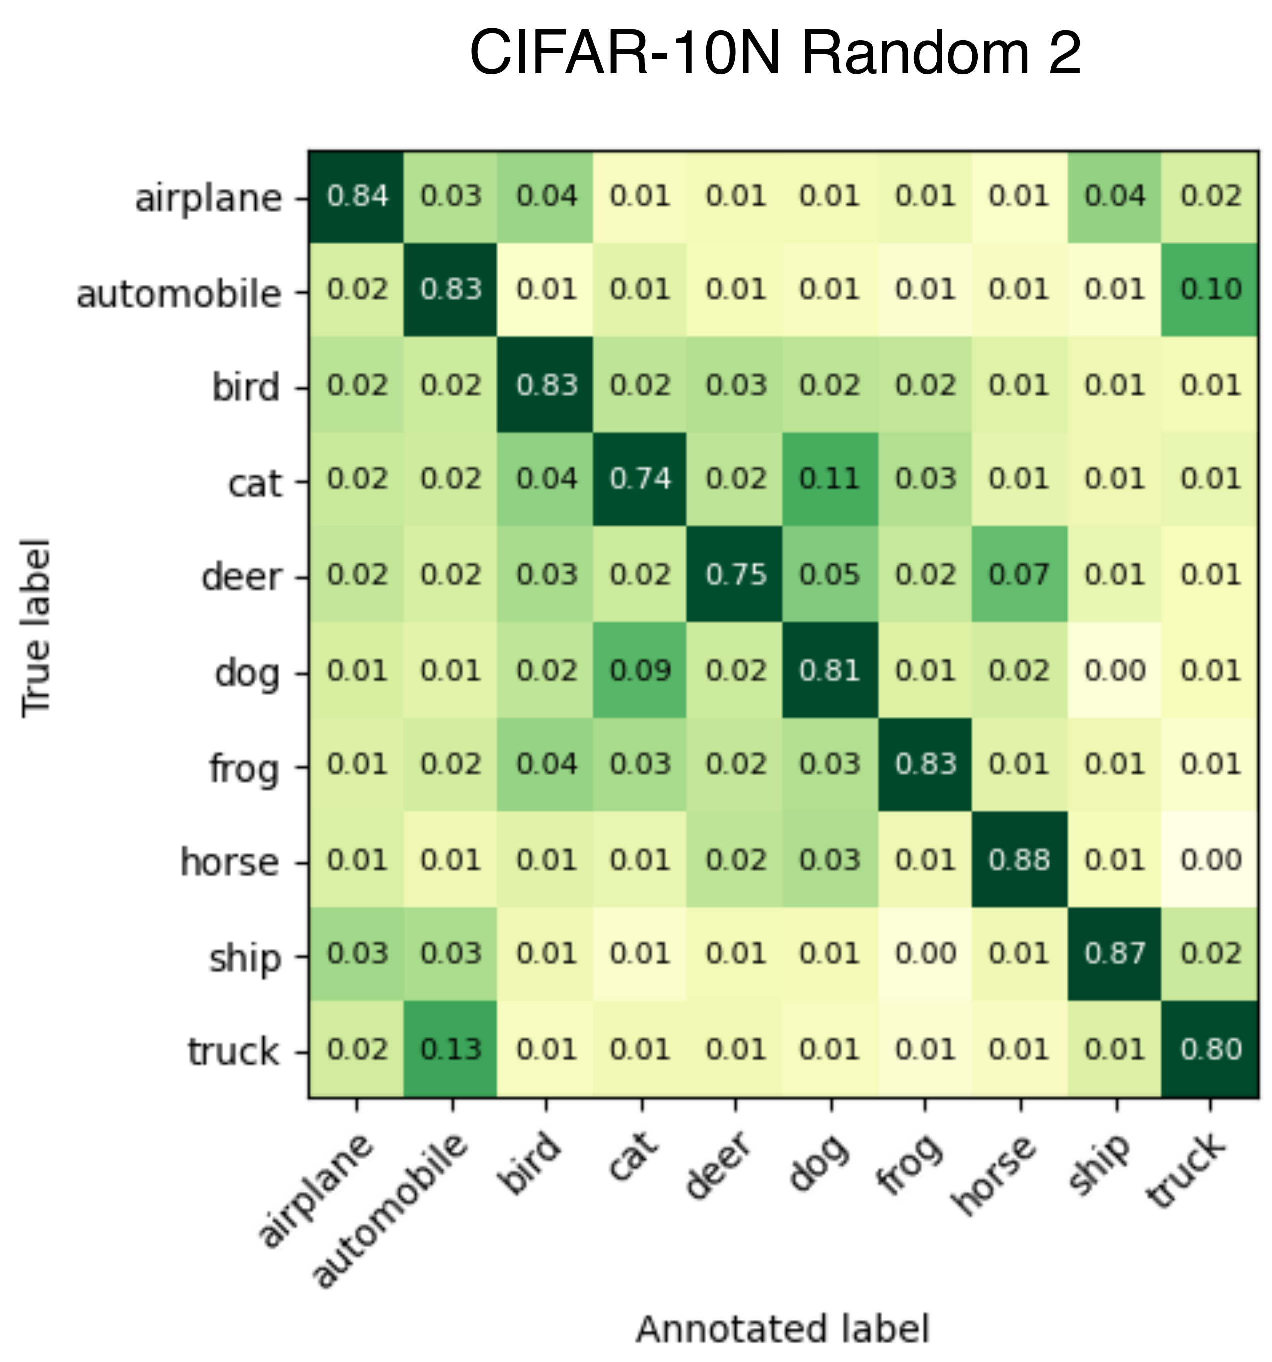

Experiments

- **Network Settings:** WideResNet-40-2 model **[10]**, SGD optimizer, cosine annealing scheduler **[11]**
- **Data preprocessing:** random horizontal flips and batch normalization
- ***AugLoss* Settings:**

Augmentation	**Loss Function**
NoAug (baseline)	CE loss (baseline)
AugMix	Focal loss
	NCE+RCE loss
	Alpha-loss

Experiments

- **Result #1:** *AugLoss* (i.e., AugMix + robust loss) appears to combat the tested settings of real-world dataset corruption, performing the best in all label noise categories

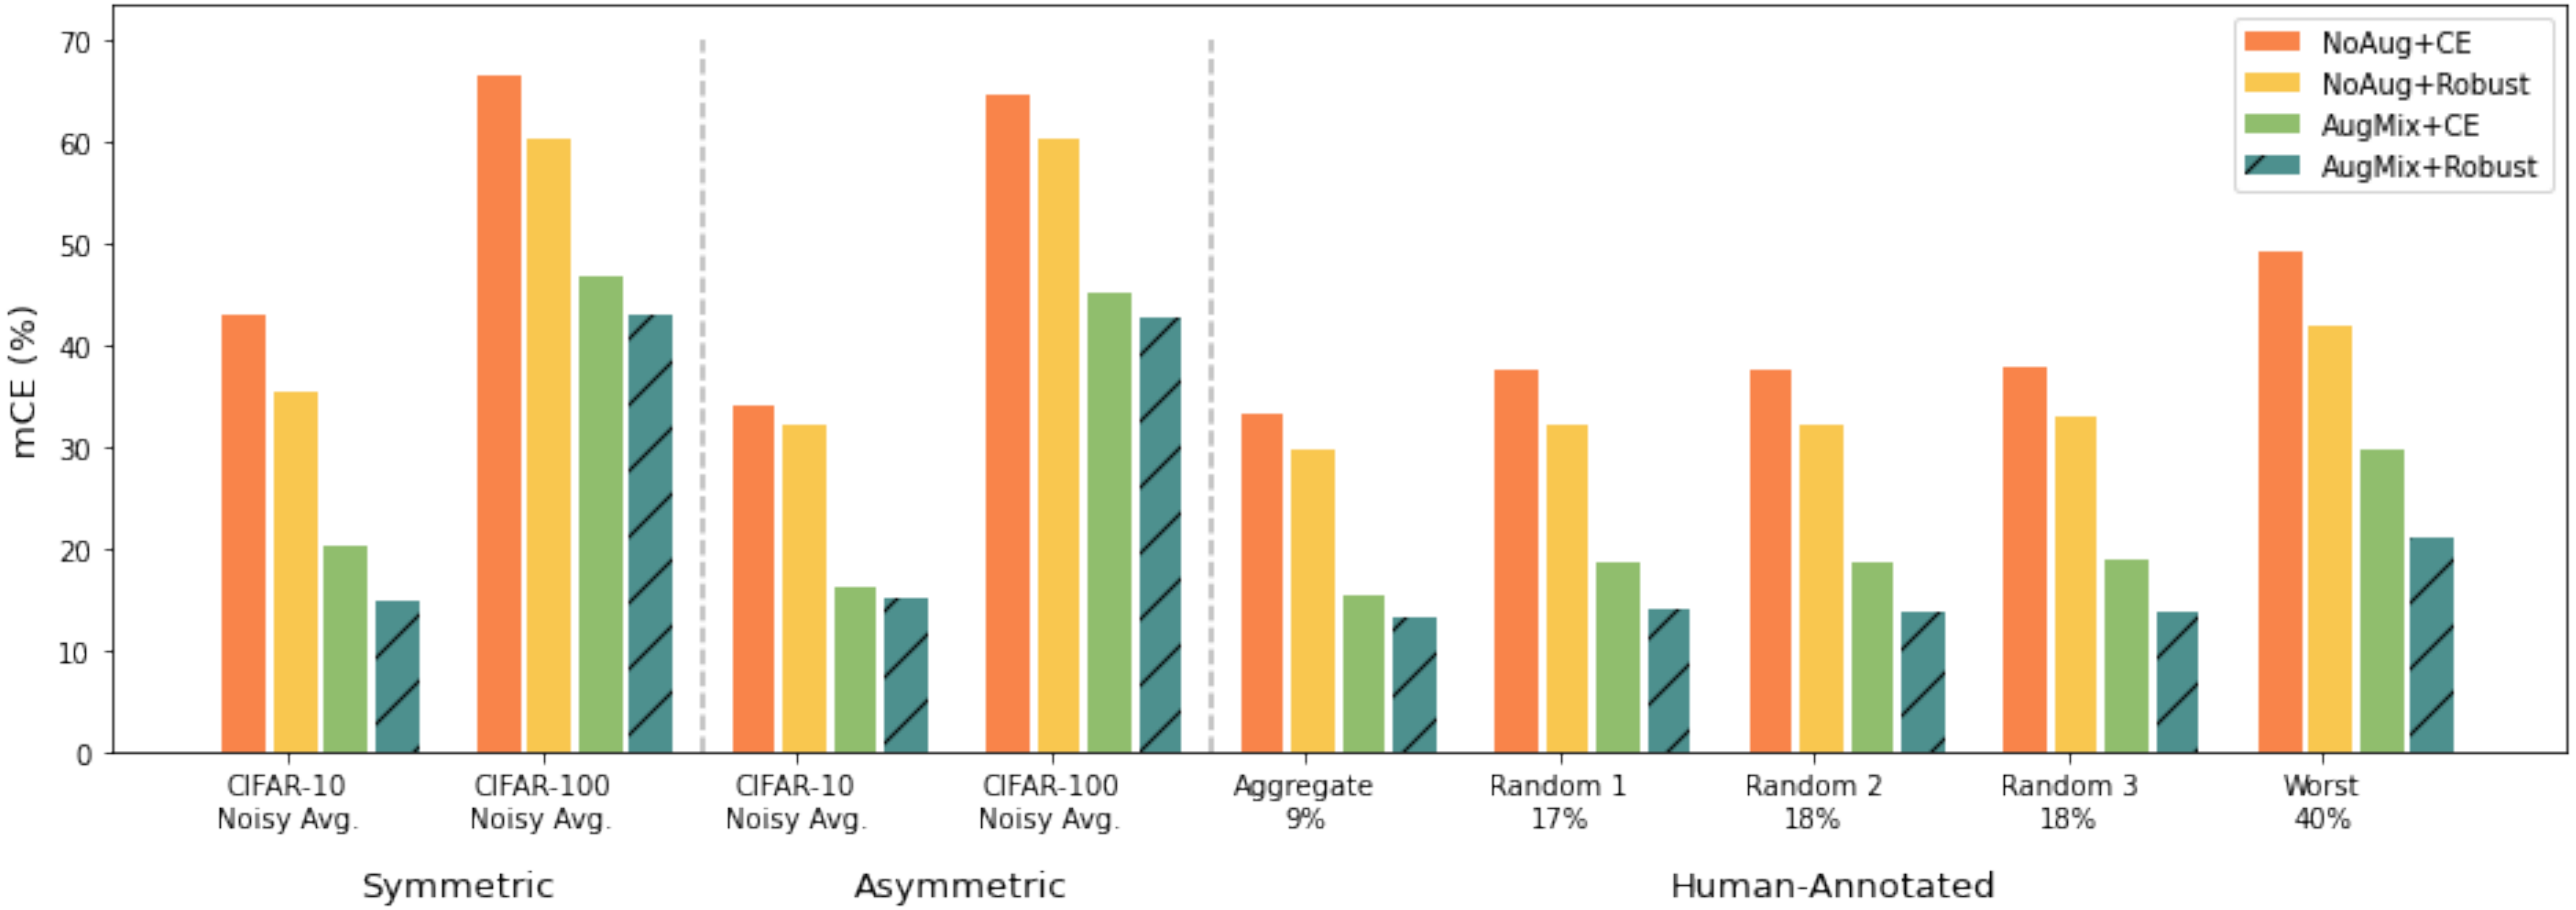

Experiments

- **Result #2:** No specific robust loss function appears to be the "universal fit" for all tested settings of dataset corruption; rather, a mixture of losses yields the best results

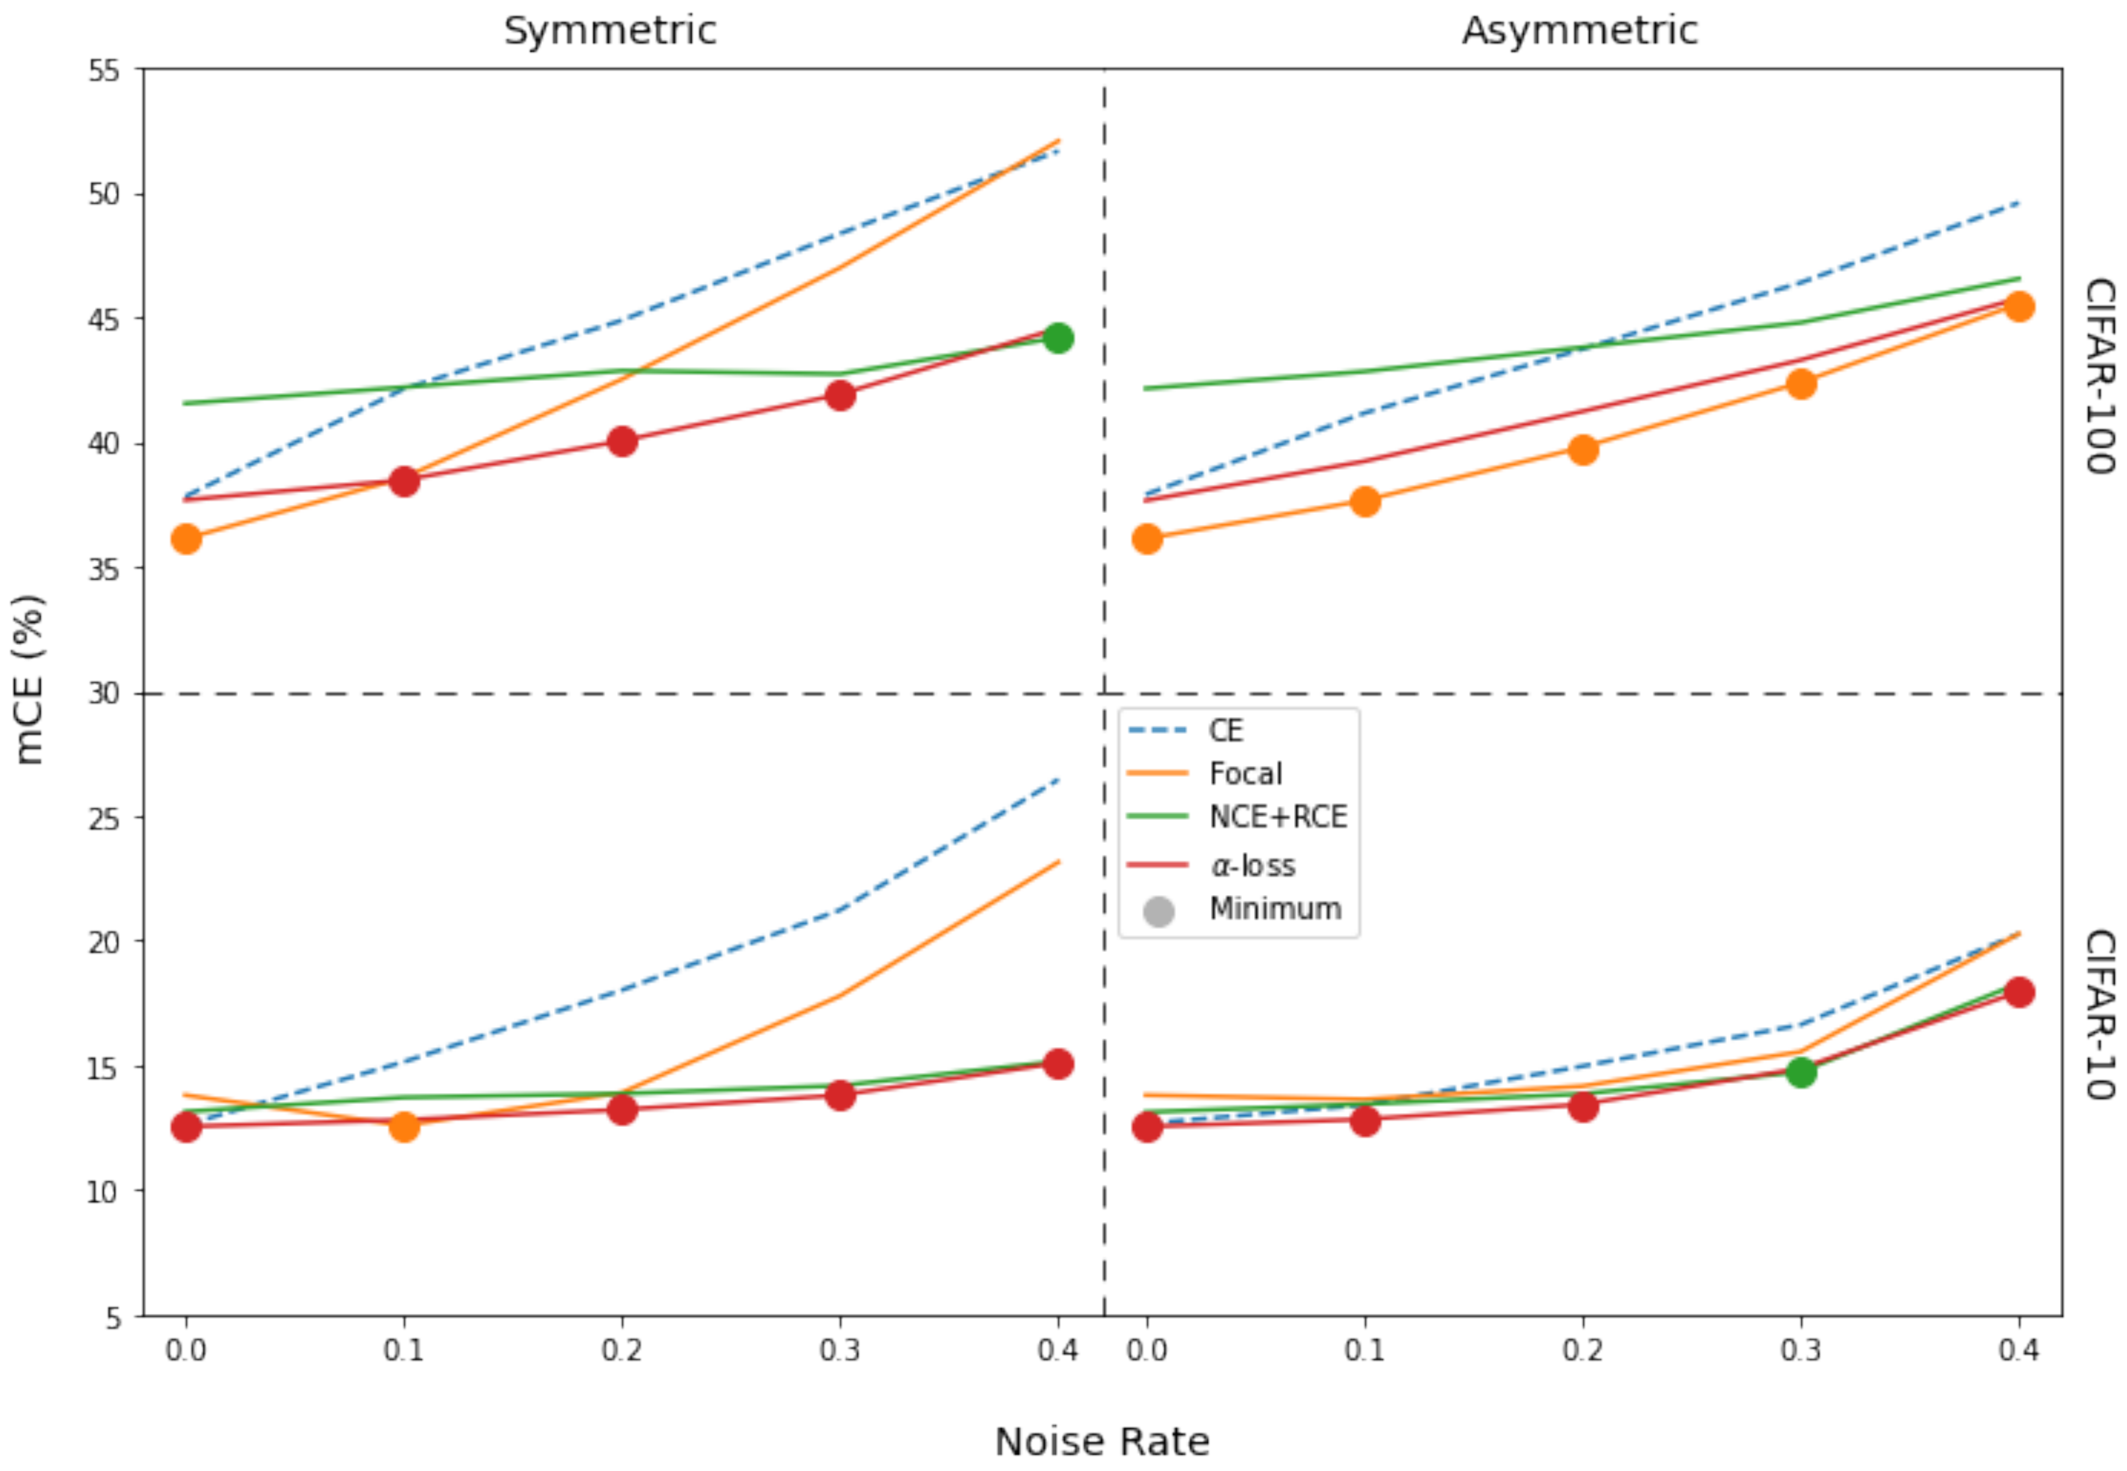

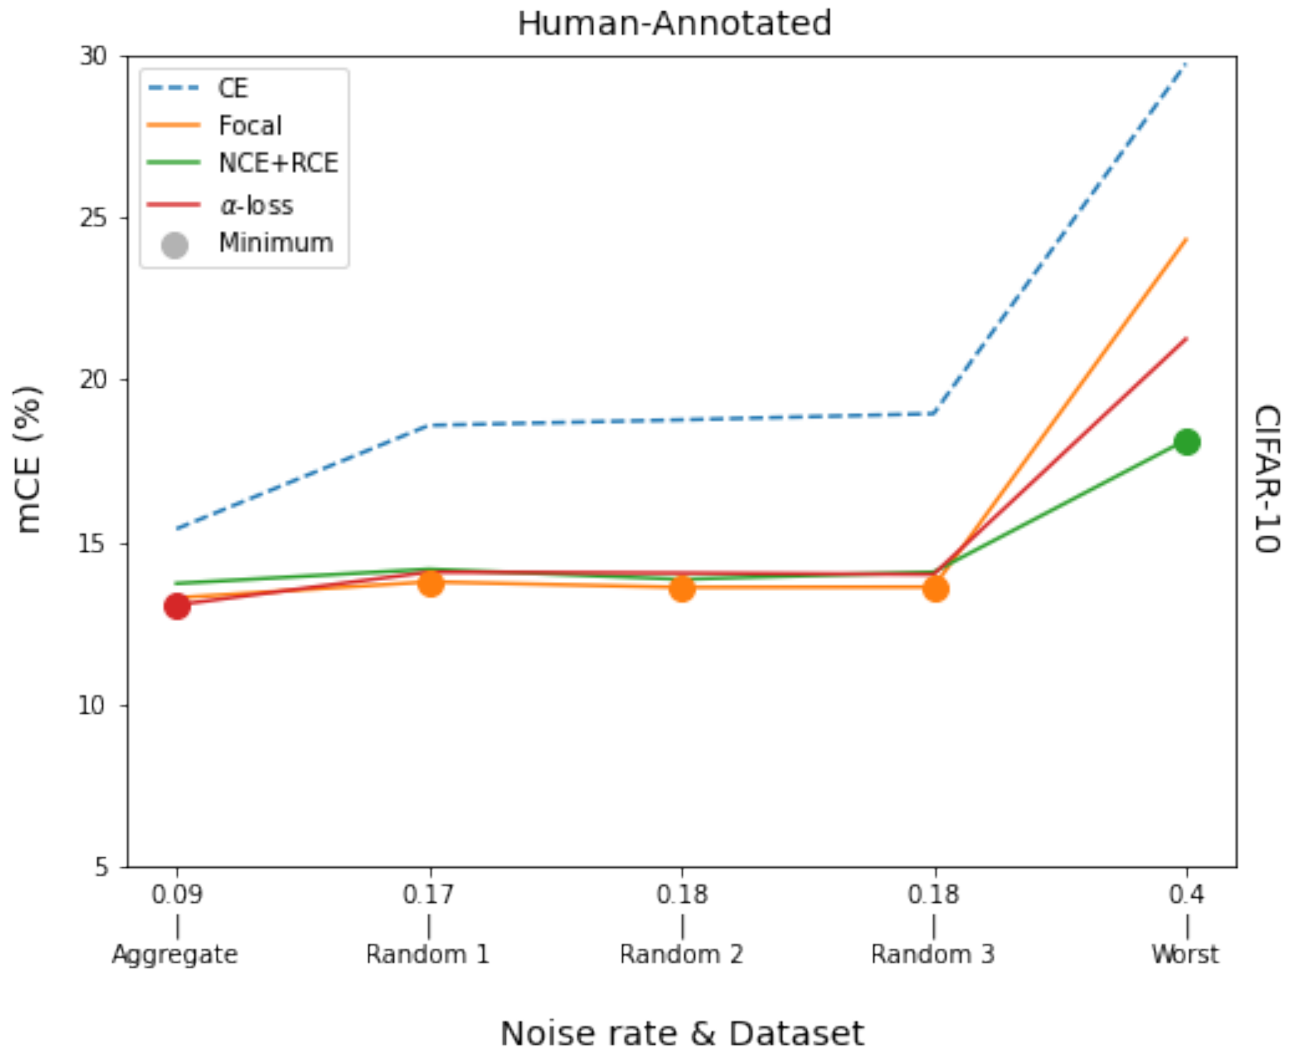

Conclusion

- **Takeaways**
  - Proposed *AugLoss*, a novel methodology combining data augmentation and robust loss functions to combat noisy labeling and test-time distribution shifts
  - Experimentally demonstrated that *AugLoss* methods can exhibit greater robustness to dataset corruption than the use of either data augmentation or robust loss alone
- **Future Work**
  - Potentially build on the efficacy of *AugLoss* by leveraging the new WILDS dataset [14] that encapsulates real-world distribution shifts

References

[1] J. Wang, C. Lan, C. Liu, Y. Ouyang, T. Qin, W. Lu, Y. Chen, W. Zeng, and P. S. Yu, "Generalizing to unseen domains: A survey on domain generalization," 2021.

[2] 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*, ser. Proceedings of Machine Learning Research, K. Chaudhuri and R. Salakhutdinov, Eds., vol. 97. PMLR, 09–15 Jun 2019, pp. 5907–5915. [Online].

[3] D. Arpit, S. Jastrzebski, N. Ballas, D. Krueger, E. Bengio, M. S. Kanwal, T. Maharaj, A. Fischer, A. Courville, Y. Bengio *et al.*, "A closer look at memorization in deep networks," in *International conference on machine learning*. PMLR, 2017, pp. 233–242.

[4] D. Hendrycks and T. Dietterich, "Benchmarking neural network robustness to common corruptions and perturbations," 2019.

[5] X. Ma, H. Huang, Y. Wang, S. Romano, S. Erfani, and J. Bailey, "Normalized loss functions for deep learning with noisy labels," 2020.

[6] 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. [Online]. Available: http://dx.doi.org/10.1109/ICCV.2017.324

[7] 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.

References

[8] 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,” *arXiv preprint arXiv:1912.02781*, 2019.

[9] 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.

[10] S. Zagoruyko and N. Komodakis,“Wide residual networks,”*arXiv preprint arXiv:1605.07146*, 2016.

[11] I. Loshchilov and F. Hutter, “SGDR: stochastic gradient descent with restarts,” CoRR, vol. abs/1608.03983, 2016. [Online]. Available: http://arxiv.org/abs/1608.03983