KIDELTA LEARNING
Scalable AI for Automated Driving

Improving robustness against common corruptions with frequency biased models

Tonmoy Saikia, Cordelia Schmid, Thomas Brox

WP 2.4
WP 2.5

Use-Case in Delta Learning

In an autonomous driving scenario, models should be able to generalize across domains. For instance, a model trained on clean images taken under daylight conditions should be **robust** to perturbations such as weather conditions like: fog, snow rain etc.

Technical Problem

Robustness to distribution shift is possibly the core challenge in deep learning. CNNs show strong performance when training and test set samples are independent and identically distributed. Unlike human vision, CNNs are severely affected even by small perturbations, for example random noise or weather distortions like snow, fog etc. (see Fig. 1). The most successful remedies to-date is well-chosen data augmentation schemes. However, data augmentation comes with robustness trade-offs, i.e., many transformations improve performance on some types of corruptions but reduce performance on clean images. In realistic scenarios, the dominant fraction of data is typically clean and not corrupted. Therefore, clean performance must not be ignored.

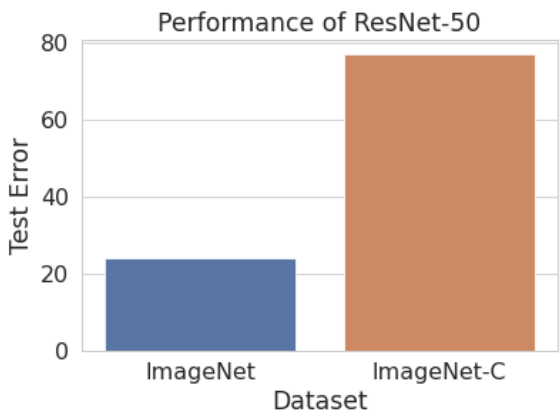

Figure 1: Adding random noise to the ImageNet test set is sufficient to almost triple classification error.

Technical Solution

We propose RoHL - Robust mixture of a HF (high-frequency) and a LF (low-frequency) expert model. To build the HF expert model, we apply TV minimization on the activations of the first convolutional layer, as well as generic augmentations that affect high-frequency components in the image. The HF expert is robust to high-frequency corruptions whereas the LF expert, based on plain contrast augmentation, is robust to low-frequency corruptions. We show that having such complementary models improves performance both on corrupted and clean images. Also compared to a standard two-member ensemble it adds robustness at no additional cost. See Fig. 2 for an overview.

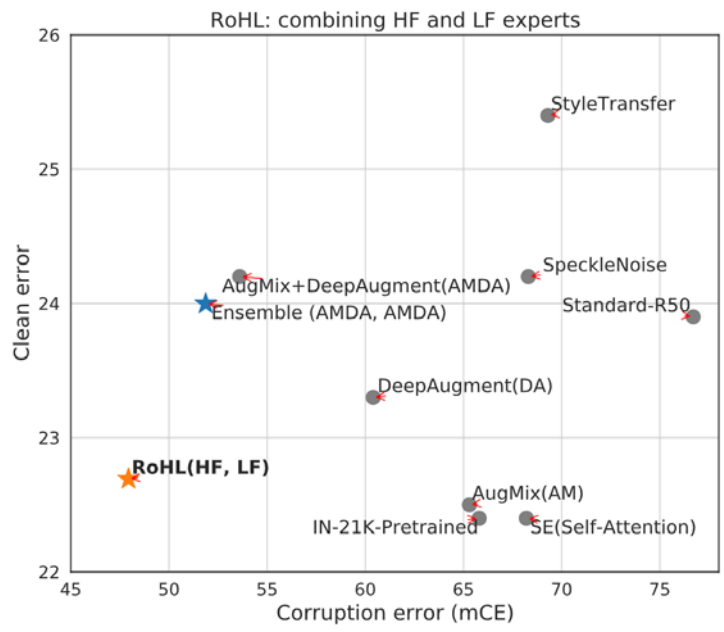

Figure 2: Improvement pf RoHL on clean and corruption errors compared to other approaches

Evaluation

We evaluated our approach on both synthetic corruptions (ImageNet-C) and real corruptions (BDD100K, DAWN – automotive datasets). We considered both object classification and detection tasks. The performance of our approach on real corruptions for the detection task is shown in in Table 1.

Pretrained Backbone	Clear AP	mAPc	Fog AP	Rain AP	Sand AP	Snow AP
Standard data augmentation	31.3	24.9	21.5	25.1	24.8	21.7
AMDA	32.4	27.2	24.9	26.2	27.6	24.8
Ensemble(AMDA, AMDA)	32.4	27.2	**25.4**	**26.2**	27.6	24.2
RoHL (AMDA$_{TV}$-ft_{Gauss},AMDA-ft_{Cont})	**32.6**	**28.8**	24.9	24.9	**28.1**	**33.4**

Table 1: Object detection performance with different ResNet50 backbones (pretrained on ImageNet) used in FasterRCNN. We report average precision (AP) scores on the "Clear" split of BDD100k and corrupted testsets in DAWN. Higher AP scores are better. mAPc denotes the mean AP over corruption types.

Environment

For more information contact:
saikiat@cs.uni-freiburg.de

Partners

External partners

KI Delta Learning is a project of the KI Familie. It was initiated and developed by the VDA Leitinitiative autonomous and connected driving and is funded by the Federal Ministry for Economic Affairs and Energy.

Supported by:
Federal Ministry for Economic Affairs and Energy
on the basis of a decision by the German Bundestag