Improving robustness against common corruptions with frequency biased models

Tonmoy Saikia, Cordelia Schmid, Thomas Brox

KI DELTA LEARNING
Scalable AI for Automated Driving

Zero-shot robustness to distribution shifts

Goal: Build a model that handles low-level distribution shifts between training and test data faithfully.
Zero-shot robustness: The model should be robust to such changes without observation of the actual distribution shift.
Avoid drop of in-distribution performance.

Typical low-level shifts:

- Noise, blur, contrast changes
- Fog, rain, sand, snow

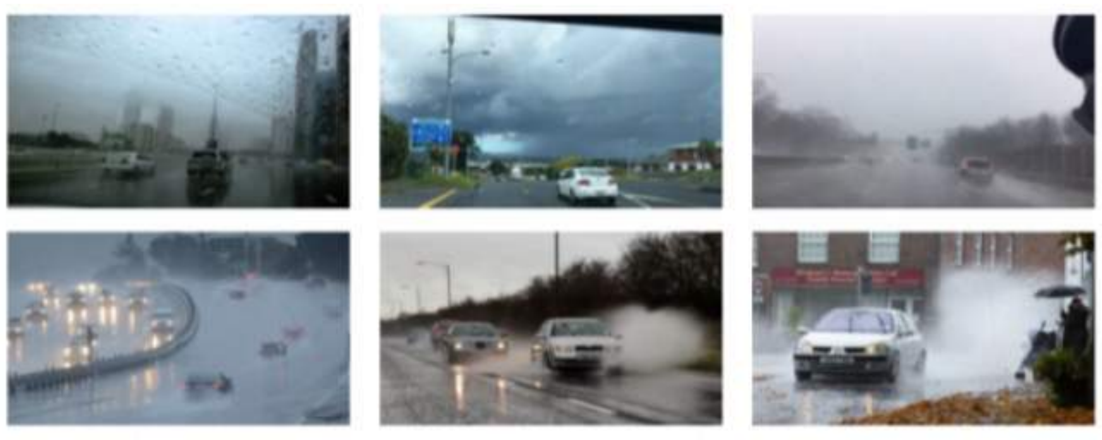

Figure 1. Test samples with rain from the DAWN dataset
© DAWN dataset (Kenk & Hassaballah)

Approach

Observation: Corruptions affecting the low-frequencies of the images (noise, blur) require different network regularization than corruptions that affect the low-frequencies (contrast, brightness).

→ Train a low-frequency and a high-frequency expert and combine them as a mini-ensemble.

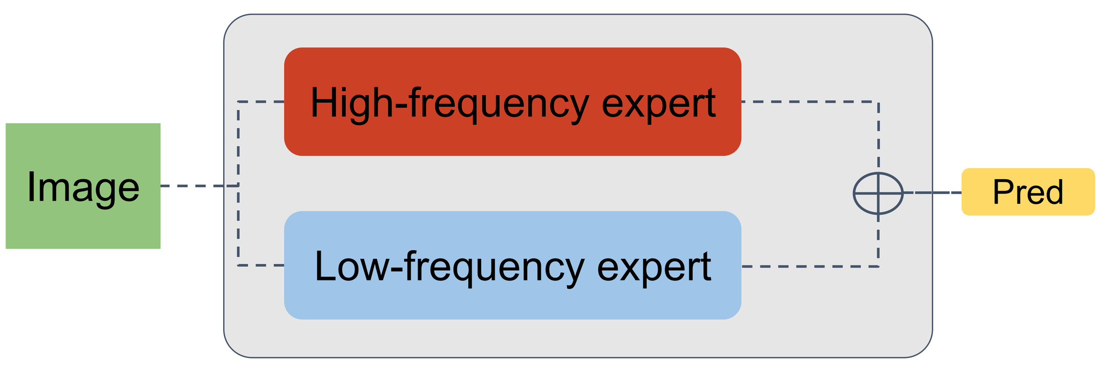

High-frequency expert:

- Trained with Gaussian noise and blur
- TV-regularization of first-layer feature map

$$\mathcal{L}(\bar{\mathbf{y}}, \mathbf{y}, \mathbf{F}) = \mathcal{L}_{CE}(\bar{\mathbf{y}}, \mathbf{y}) + \lambda \sum_c \mathcal{L}_{TV}(\mathbf{F}_c)$$

Low-frequency expert:

- Trained with contrast augmentations

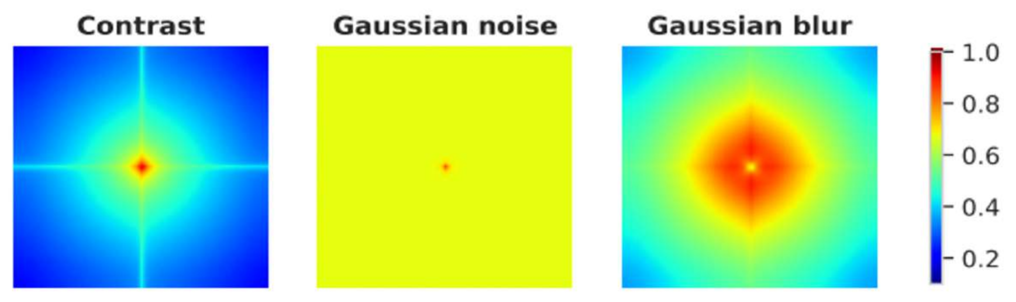

Figure 2. Fourier spectra of three corruptions used for augmentation
© University of Freiburg

Code

[o] Vision
COMPUTER VISION University of Freiburg

Results on ImageNet

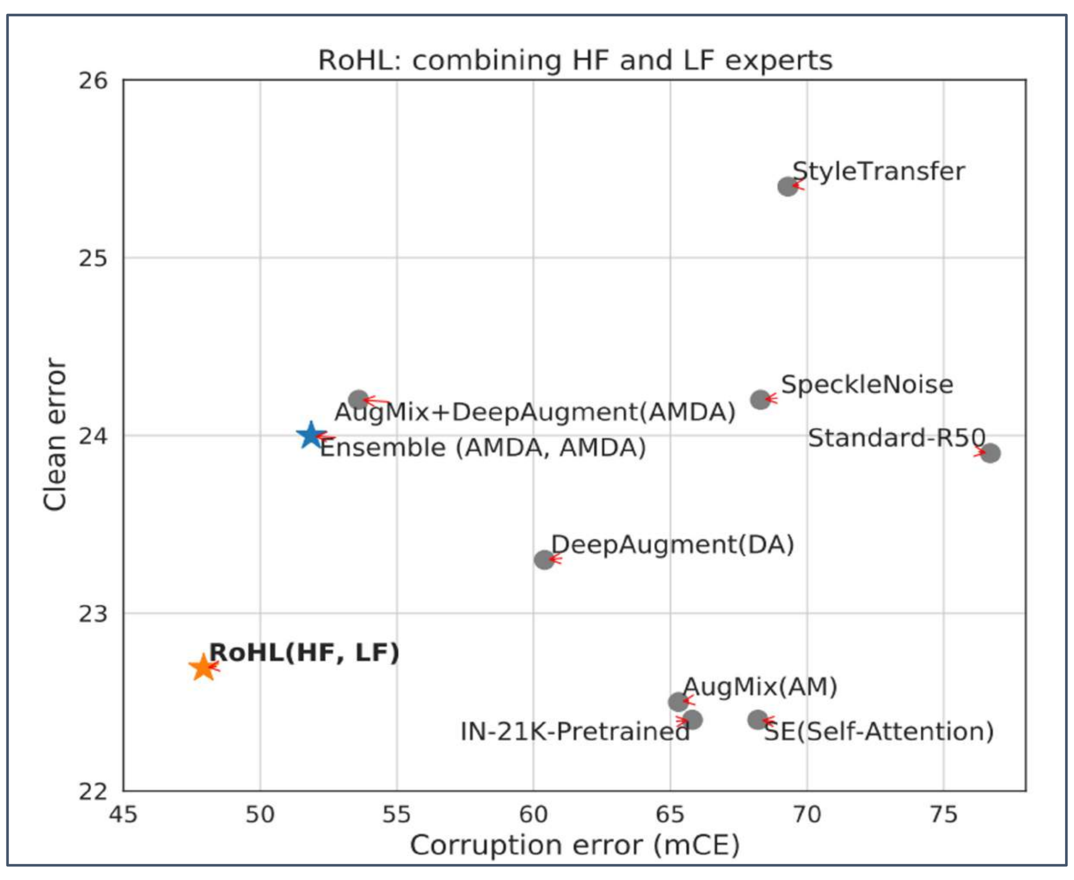

- Best corruption error of all existing ResNet-50 approaches.
- Excellent trade-off between in-distribution performance (clean error) and out-of-distribution performance (corruption error).

Results on Real-world corruptions

The DAWN dataset comprises real-world distribution shifts not seen during training. Tested by changing the backbone of Faster-RCNN.

Pretrained Backbone	Clear AP	mAPc	Fog	Rain	Sand	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**

- Particularly good with Sand and Snow

Figure 3. Test samples with sand and snow from the DAWN dataset
© DAWN dataset (Kenk & Hassaballah)

Compatible with adaptation of batch statistics (Schneider et al.)

Model	ImageNet-C mCE		DAWN-cls mCE	
	base	adapt	base	adapt
Standard	76.7	62.2	23.5	16.8
AMDA	53.6	45.4	16.4	13.6
Ensemble(AMDA, AMDA)	51.9	44.7	16.2	13.5
RoHL (AMDA_{TV}-ft_{Gauss},AMDA-ft_{Cont})	**47.9**	**41.2**	**14.5**	**12.4**

References:

- Tonmoy Saikia, Cordelia Schmid, Thomas Brox. Improving robustness against common corruptions with frequency biased models, IEEE International Conference on Computer Vision (ICCV), 2021.
- Dan Hendrycks, Norman Mu, Ekin D. Cubuk, Barret Zoph, Justin Gilmer, and Balaji Lakshminarayanan. AugMix: A simple data processing method to improve robustness and uncertainty. ICLR, 2020.
- Dan Hendrycks, Steven Basart, Norman Mu, Saurav Kadavath, Frank Wang, Evan Dorundo, Rahul Desai, Tyler Zhu, Samyak Parajuli, Mike Guo, Dawn Song, Jacob Steinhardt, and Justin Gilmer. The many faces of robustness: A critical analysis of out-of-distribution generalization. arXiv 2006.16241, 2020.
- Steffen Schneider, Evgenia Rusak, Luisa Eck, Oliver Bringmann, Wieland Brendel, and Matthias Bethge. Improving robustness against common corruptions by covariate shift adaptation. NeurIPS, 2020.

Environment

Partners

External partners

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

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 Climate Action.

www.ki-deltalearning.de @KI_Familie KI Familie

KI FAMILIE

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