Course Project Proposal

Robust Texture Descriptors

Background

In [1] Geirhos et al. found that CNNs which were pretrained on ImageNet seem to have a “texture-bias”. This means when given an image with a texture-shape cue conflict as shown in figure 1 (taken from [1]), such as a cat which was given elephant skin through style transfer, the ImageNet-trained CNNs showed a tendency to classify for texture rather than shape, i.e. classified the image in fig. 1c as “elephant” rather than “cat”.
Furthermore, it was found that CNNs, unlike humans, performed surprisingly weak in classifying the same images when given silhouette or edge-extracted images as shown in figure 2 (taken from [1]).
The authors then train CNNs by providing images such as the cat in fig. 1b with many different textures, which can be regarded as an extensive way of data augmentation. It leads to CNNs which are essentially texture-invariant and strong shape classifiers, robust against local disturbances and variations such as noise.

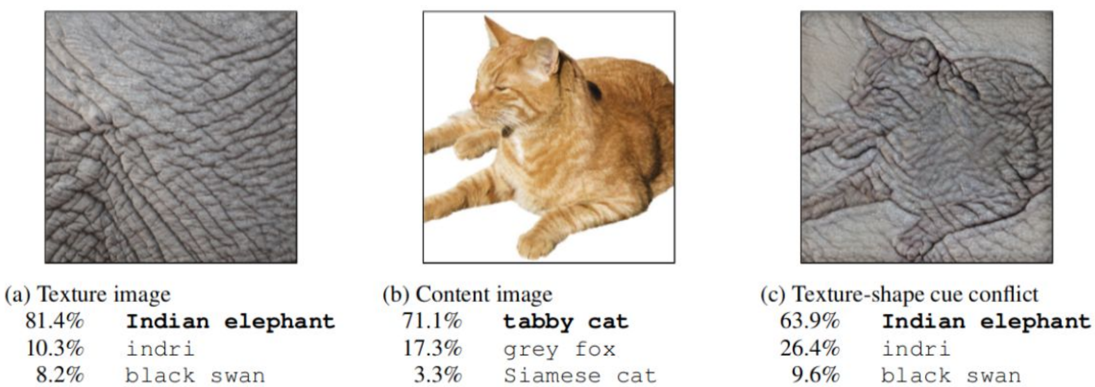

Figure 1: Classification of a standard ResNet-50 of **(a)** a texture image (elephant skin: only texture cues); **(b)** a normal image of a cat (with both shape and texture cues), and **(c)** an image with a texture-shape cue conflict, generated by style transfer between the first two images.

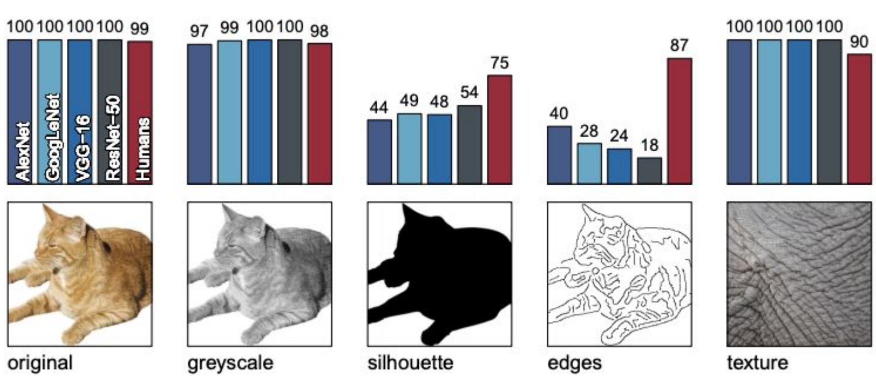

Figure 2: Accuracies and example stimuli for five different experiments without cue conflict.

Course Project Proposal

Project Description

The task is to use a similar approach as the authors of [1], but instead of making the CNN texture invariant, with an aim to produce a robust, all-purpose texture descriptor by turning the CNN shape invariant instead. This should be tried by using a texture dataset such as DTD [2] and style transfer as suggested in [1], to map textures present in the DTD dataset onto a magnitude of different underlying shapes provided in the ImageNet dataset. By classifying for texture instead of shape, the idea is that we train the CNN as a feature extractor that is robust in classifying texture. The trained CNN should then be evaluated and compared to the performances of other state-of-the-art methods in the classification of texture sensitive datasets, e.g. in biomedical applications such as oral cancer detection [3], or lymphoma classification [4].

The project work should include

- Preparation of the project plan and distribution of the tasks within the team.
- A survey of the relevant literature.
- Implementation of the selected methods in a common environment (e.g., Matlab or Python, preferably in pytorch).
- Quantitative evaluation of the selected methods on the provided data.
- Writing of the project report.

References

[1] R. Geirhos et al. "ImageNet-trained CNNs are biased towards texture; increasing shape bias improves accuracy and robustness", ICLR 2019
Code available at: https://github.com/rgeirhos/texture-vs-shape
[2] M.Cimpoi, S. Maji, I. Kokkinos, S. Mohamed, A. Vedaldi. "Describing Textures in the Wild", CVPR 2014.
[3] H. Wieslander, G. Forslid, E. Bengtsson, C. Wählby, J.M. Hirsch, C.R. Stark, S.K. Sadanandan "Deep Convolutional Neural Networks For Detecting Cellular Changes Due To Malignancy", ICCV 2014.
[4] N.V. Orlov et al. "Automatic Classification of Lymphoma Images with Transform-Based Global Features", IEEE Trans Inf Technol Biomed 2010

Contact

Elisabeth Wetzer
elisabeth.wetzer@it.uu.se

Joakim Lindblad
joakim@cb.uu.se

Nataša Sladoje
natasa.sladoje@it.uu.se

Centre for Image Analysis, Department of Information Technology, Uppsala University