Psychology

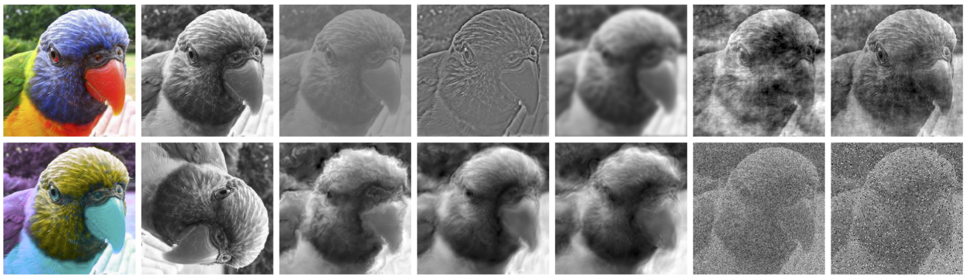

Figure 1.11: Human-DNN-robustness comparison: stimulus images. From left to right, image manipulations in the top row are: undistorted, greyscale, low contrast, high-pass, low-pass (blurring), phase noise, power equalisation. Bottom row: opponent colour, rotation, Eidolon I, II and III, additive uniform noise, salt-and-pepper noise. Typically human observers are more robust when the signal-to-noise ratio decreases. Additionally, we find progressively diverging patterns of classification errors between humans and DNNs with weaker signals.

This project investigates human perception, combining psychophysical experiments and computational modeling. We have five research foci:

First, we develop methods to gain more information from psychophysical data, linking traditional methods with machine learning approaches [80, 246]. We characterized serial dependencies in behavioral responses and introduced methods to correct for them. Our software package to perform Bayesian inference for the psychometric function for non-stationary data is freely available and widely used [81]. In addition, we investigate reliable supra-threshold psychophysical paradigms which are more intuitive and require less training and are thus more likely to yield reliable crowd-sourcing data [54, 63].

Our second focus is the development of an image-based model of spatial vision. We integrated the large psychophysical literature on detection and discrimination experiments and proposed a model based on maximum-likelihood decoding of a population of model neurons predicting the most important spatial vision data sets simultaneously, using a single set of parameters [56]. Remarkably, the model generalises well to data it was not trained upon (natural images), produces a highly sparse code, and is useful in applications (in perceptual visual quality metrics). In the future, we plan to assay whether the model could be useful as a preprocessing step for deep neural networks (DNNs).

Third, we are exploring similarities and differences of DNNs and the human visual system. One line of work uses DNNs as generative models for texture synthesis, casting doubt on the popular notion that the peripheral visual system's internal representations are texture-like [55]. In a second line of work, we compare the robustness of humans and current convolutional DNNs on object recognition under various types of image degradation, finding the human visual system (still) to be much more robust [109]. These differences cannot be overcome by training on distorted images (i.e., data augmentation): While DNNs cope well with the exact distortion they were trained on, they still show a strong generalisation failure towards unseen distortions.

Fourth, we investigate perception of causality, focusing on the arrow-of-time that we previously studied theoretically [234]. Preliminary data suggest that human observers can discriminite forward and backward played movies of autoregressive (AR) motion with non-Gaussian additive independent noise, i.e., they appear sensitive to the subtle temporal dependencies of the residuals of the AR-motion. Currently we are testing how long the motion sequences have to be for successful arrow-of-time discrimination and how well observers can generalise from one type of non-Gaussian noise to another.

Another project pertaining to the wider field of psychology is described in the project report on social networks, dealing with the optimization of spaced repetition algorithms for learning [17].

More information: https://ei.is.mpg.de/project/psychology