

# Can the early human visual system compete with Deep Neural Networks?

Samuel Dodge and Lina Karam  
Arizona State University  
{sfdodge, karam}@asu.edu

## Abstract

We study and compare the human visual system and state-of-the-art deep neural networks on classification of distorted images. Different from previous works, we limit the display time to 100ms to test only the early mechanisms of the human visual system, without allowing time for any eye movements or other higher level processes. Our findings show that the human visual system still outperforms modern deep neural networks under blurry and noisy images. These findings motivate future research into developing more robust deep networks.

## 1. Introduction

It has been shown that state-of-the-art deep neural networks (DNNs) often achieve better performance compared with human subjects on large scale classification tasks [18]. Given this, we might conclude that the DNNs do a better job of representing and organizing visual data. But does the success of DNNs carry over to other more difficult tasks?

Recently it has been discovered that deep neural networks perform poorly in the presence of distortions such as blur and noise [6, 17]. Blur removes high frequency information and likewise, noise injects high frequency information. Current deep networks seem to experience difficulty reasoning in the presence of high levels of such distortions.

Do humans have a similar trouble with distorted images? By studying the human visual system (HVS), we can perhaps gain insight into how to build DNN models that are more robust to distortions. If human performance on distorted images is better than DNNs, then this exposes a vulnerability in the DNN's representation of visual data. But if human performance is also poor, then recognition under distortions may be inherently difficult.

Previous studies have tested human capability for recognition under noise and blur [22, 5], finding that humans have some robustness with respect to these distortions. In this work, we wish to compare human ability and DNN ability on a common task. Recent works perform similar experiments with distorted images [7, 10]. However, in this work

we wish to test the *early human vision system* by limiting the stimuli display time to 100ms. Within 100ms, there is no time for eye movements [1], thus the human visual system is limited to more global “gist” representations [14]. Can the human visual system still recognize distorted images only by the gist? By contrast, the experiments in [7] allow the subject to view the image for unlimited duration. This allows the subject to analyze more local information that can help classify the stimuli.

### 1.1. Related Works

Human performance on distorted stimuli has been extensively studied. Torralba *et al.* [22] showed that humans are able to recognize very low resolution images. Similarly, studies on face images show that the human visual system can perform well in the presence of blur [2] and noise [5].

There are also several works that study deep neural network performance on distorted data. Dodge and Karam [6] studied several different types of common image distortions and found that noise and blur have the largest effect on the performance of DNNs on the ImageNet dataset. Rodner *et al.* show similar findings on smaller fine grained datasets [17].

Comparing human and machine vision performance has also been studied in the past. Borji and Itti [3] compare 14 different computer vision models on several datasets and compare with human performance. However the study does not consider modern neural networks, which greatly outperform older vision models.

Fleuret *et al.* compare human and machine vision performance on synthetic visual reasoning tasks [9]. These tasks are designed such that there must be some sort of reasoning, instead of pure pattern recognition. For many of these tasks humans outperform artificial vision systems. A followup study showed that, for some of these synthetic problems, state-of-the-art neural networks surprisingly achieve accuracy equivalent with random chance [20].

Parikh [15] studies human and artificial vision system performance on jumbled images. These jumbled images are formed by randomly permuting blocks of an image. On these jumbled images human classification performance is

degraded to near the performance of a bag-of-words based classifier. While the jumbled images may give some insight into the human visual system, jumbled images are not typical visual stimuli.

Kheradpisheh *et al.* [13] compare human and DNN performance on images of objects with arbitrary backgrounds and rotations. The highest performing DNNs match human performance. This is consistent with other studies (e.g., [18]) that show that DNN classification performance is at-par with or superior to human performance.

To further evaluate human vs. DNNs for classification we design our experiments to test highly distorted images. We extend on the work in [7], where Amazon Mechanical Turk testing is used to compare human and machine performance on a subset of the ImageNet dataset with added distortions. In this work there are two primary differences. First we use human subjects in a controlled lab setting. This is in contrast to Amazon Mechanical Turk studies where there is no mechanism to control viewing distance, screen brightness, etc. Secondly, we use a fixed display time instead of free viewing, which lets us analyze the accuracy of the HVS without allowing for higher level processes like eye movements.

A concurrent independent study [10] compares human and deep learning performance on distorted images. However, the study does not fine tune networks on distorted images. This gives an unfair advantage to the human subjects, which may have previously seen distorted images. Additionally, we chose a 100ms display time which is less than the 200ms display time in [10]. This ensures that we are testing early “gist”-based processes of the visual system.

## 2. Methods

In this section we first introduce our dataset and image distortions. Then, we describe our experimental setup to test classification under distortion by both human subjects and deep neural networks.

### 2.1. Dataset

In this paper we wish to test the early vision system. Thus we chose a simple coarse-grained dataset, instead of a more difficult fine-grained dataset (as in [7]). We select 8 classes from the Caltech101 dataset [8]: Butterfly, Crocodile, Dolphin, Elephant, Flamingo, Leopard, and Llama. This dataset was deliberately chosen to be “easy” such that under no distortion, neither the human subjects nor the deep neural networks would have any difficulty in classification.

We consider two types of distortion: random additive Gaussian noise and Gaussian blur. For Gaussian noise we consider noise with a standard deviation ranging from 0 (clean images) to 200 (highly corrupted images). Similarly, for blur we consider Gaussian blur with a kernel of standard

deviation that ranges from 0 to 10. Figure 1 shows example stimuli of the 8 classes for 3 difference levels of each distortion.

Our dataset consists of 200 clean training images (25 per class), and 40 clean validation images (5 per class). We use 80 unique images for testing blur at 5 possible levels for a total of 400 images. Similarly there are 80 unique images for testing noise at 5 levels for a total of 400 images. We additionally randomly distort the training images as explained in Section 2.3. The same training, testing, and validation split was used for each human trial, as well as for the deep neural networks.

### 2.2. Human Experiments

The goal of our human experiments is to test the ability of the human visual system to identify images using 100ms “gist”-level information. Other studies have shown that humans can accurately recognize images at display times below 100ms [12, 16]. We choose 100ms because this is the same display time in the experiments in [4] which are used to correlate human and DNN neural responses.

We design our experiment to mimic the training, validation, and testing stages used for training and evaluating DNNs.

As in [7] we first allow the subjects to freely view the training images. The subject must view images in all of the categories before allowed to continue with the experiment. This allows the subjects to familiarize themselves with the image categories. This training stage is analogous to the training stage used for deep neural networks.

Before we test the subject on distorted images, we first test clean images in a validation stage. This stage is to ensure that the subjects are able to correctly classify clean images. As in the rest of our experiments a central fixation cross is first shown for 500 milliseconds, followed by the image stimuli for 100 milliseconds, and finally a second fixation cross for 500 milliseconds (Figure 2). Next, a choice screen allows the subject to choose the most appropriate of the 8 categories and continue to the next image. This is a forced-choice response, so the subject must choose a class, even if the subject is not perfectly clear as to the correct class.

In the next stage of the experiment the subject is asked to classify distorted images. The experiment proceeds exactly as in the validation stage, except that the images are now distorted. For each image we wish to test the maximum distortion level at which the subject can make a correct decision. Testing every possible distortion level for every image can accomplish this but has practical limitations. First, with many distortion levels, the total duration of the experiment becomes prohibitive. Secondly, multiple exposure to the same image can induce a memory effect that could help subjects identify highly distorted images that they would

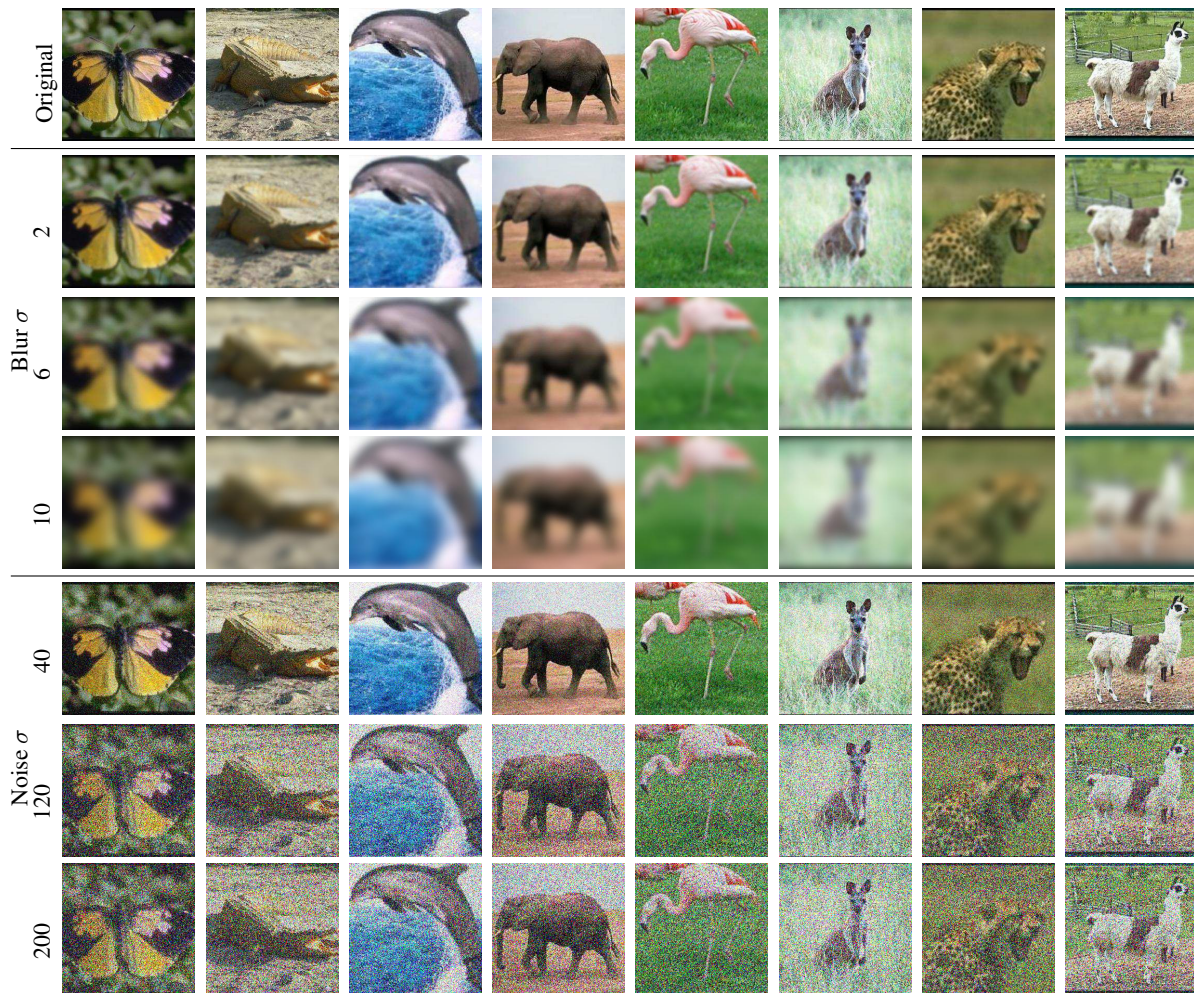


Figure 1. **Example image stimuli.** Our image categories consist of “coarse” classes that are easily recognizable under no distortion. In this figure we show three levels of distortions among the 5 levels used in the experiments.




Figure 2. **Timing of validation and testing stages.** The subject is first shown a central fixation marker for 500ms followed by the image stimuli for 100ms, another fixation marker for 500ms and finally a choice screen for the subject to input the class estimate.

otherwise not be able to identify.

Instead we follow the procedure in [7], which makes the assumption that subjects can correctly classify an image with a certain distortion level if the subject has already correctly classified an image with a higher distortion level. Thus we begin with each stimuli image at the maximum distortion level. If the subject incorrectly identifies the image, then the distortion level is reduced and the image is randomly shown later. If the subject correctly classifies an image at a particular distortion level, then all of the lower distortion levels are assumed to be correctly classified and the lower distortion levels are not tested.

We recruit 8 subjects to participate in our experiment. All of the subjects were tested for vision and color blindness before beginning the experiment.

### 2.3. Deep Neural Networks

Deep networks are the current state-of-the-art approaches for many image recognition tasks. These net-






Figure 3. **Human experiment training screen.** During the training stage, the subject must view example images from each category before continuing with the experiment.

works consist of layers of convolutional filtering, pooling, and nonlinear operations. The parameters of the layers can be learned by fitting to a training set using gradient descent based optimization algorithms.

We consider three classification models: VGG16 [19], GoogLeNet [21], and ResNet50 [11]. The VGG16 network is a popular architecture which consists of 16 layers and uses convolution layers with small 3x3 filters. The GoogLeNet architecture introduces blocks that perform parallel convolution with filters of varying size. ResNets employ skip connections, which aids in training and can yield a more accurate network.

Each network has been pre-trained on the ImageNet dataset [18]. We perform fine-tuning to adapt the model to our new dataset. Specifically, we replace the last fully connected layer with a new fully connected layer with 8 units (corresponding to the 8 classes in the dataset). The learning rate of the new layer is 10x the learning rate of the pre-trained layers. The network is fine-tuned using stochastic gradient descent with momentum. We stop learning when the performance on the validation set plateaus.

We test two scenarios for training. In the first scenario, the networks are fine-tuned on clean images as previously described. In the second scenario, the networks are fine-tuned on a mixture of clean images and distorted images. We refer to these networks as “distortion-tuned”. During each mini-batch, half of the training images remain undistorted and the other half are distorted with a random level uniformly chosen from the minimum distortion level to the maximum distortion level. This procedure is identical to that in [7], and ensures that the distortion-tuned networks can perform well at all levels of distortions, as well as for

clean images.

We test using the same procedure as in the human experiments. The network is first tested on the highest distortion level, and if the prediction is correct the network is assumed to correctly classify the same image at all lower levels of distortions. If the prediction is incorrect, then the network is tested on a lower level of distortion. This process continues until the network predicts the correct class, or the level of distortion becomes 0 (i.e., a clean image).

### 3. Results

On the validation portion of the human test, subjects scored an average accuracy of 99.3%. This shows that the subjects are diligently performing the experiment, and that the subjects are able to classify the images correctly, even at the display time of 100ms.

Figure 4 shows the comparison between human accuracy and machine accuracy for our experiments. For clean images human accuracy and DNN accuracy are nearly identical. When distortion is added, human accuracy greatly exceeds that of neural networks trained on clean images. When networks are fine-tuned on distorted data, the network performance approaches human performance for most levels of distortion. However at the largest levels of distortion, the human performance still exceeds that of the distortion-tuned deep networks.

The network architecture can affect how resilient the network is to distortion, even when the network is distortion-tuned. For our dataset, the distortion-tuned ResNet50 model performs the best on noisy images, but is not the best performing on blurry images. For blurry images the distortion-tuned VGG16 model achieves closest to human perfor-




Figure 4. **Accuracy comparison between DNNs and human subjects for distorted stimuli.** Human subjects greatly outperform DNNs trained on clean data. When networks are trained on data from the respective distortion, the performance gap decreases.

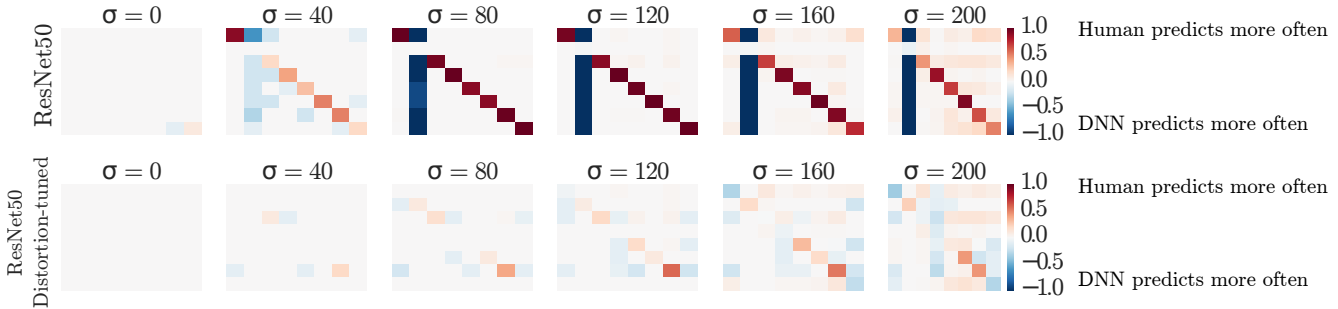


Figure 5. **Confusion matrix difference between humans and ResNet50 for noise.** Columns of a confusion matrix correspond to the actual class, and rows correspond to the predicted class. Rows are normalized to sum to one. This figure shows the *difference* between the human confusion matrix and that of the ResNet50 model. Red regions denote pairs that humans predict more often, and blue regions denote pairs that networks predict more often.

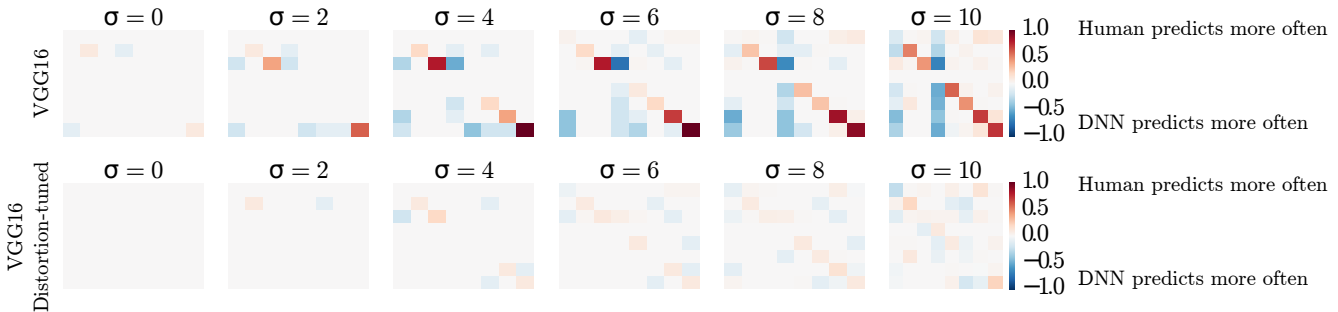


Figure 6. **Confusion matrix difference between humans and ResNet50 for blur.** Columns of a confusion matrix correspond to the actual class, and rows correspond to the predicted class. Rows are normalized to sum to one. This figure shows the *difference* between the human confusion matrix and that of the VGG16 model. Red regions denote pairs that humans predict more often, and blue regions denote pairs that networks predict more often.

mance.

An analysis of confusion patterns can yield insight into the behavior of the DNNs. We compute the difference between the confusion matrices of human subjects and the best performing DNN model. This tells us which categories are confused by the DNN relative to the human and vice-versa. Figure 5 shows this for noise using ResNet50, and Figure 6 shows this for blur using VGG16. For a DNN that has not seen distorted data during training, at high distortions the network tends to predict a single class regardless of the input. The network fine-tuned on distortions shows a confusion matrix that is more similar to humans.

We also show examples of images that are “difficult” or “easy” for humans or the distortion-tuned ResNet50 model (Figure 7). We consider an image to be correctly classified by the human subjects if 90% or greater of the subjects selected the correct class. Some images are easy for both the neural network and human subjects. Others are consistently misclassified by both. Finally there are some images correctly classified by humans, but misclassified by DNNs. Likewise some images classified correctly by DNNs are classified incorrectly by the human subjects.

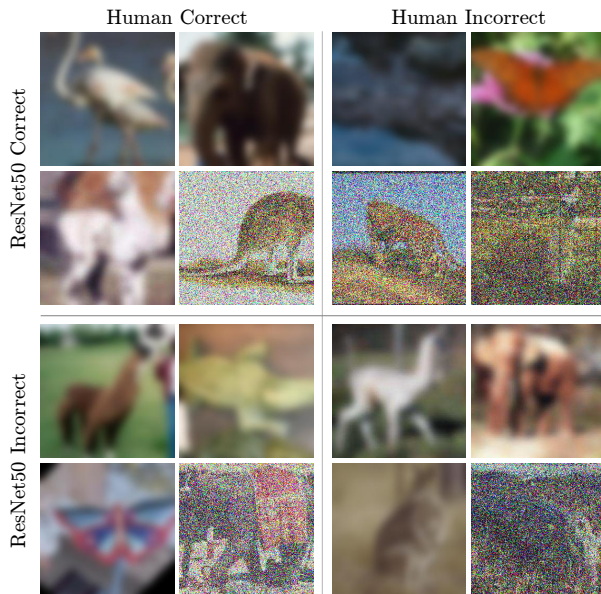


Figure 7. Examples of correct and incorrect classifications of human subjects and distortion-tuned ResNet50 models.

#### 4. Discussion and Conclusion

We performed a set of human experiments for classifying distorted images under a 100ms time constraint. We find that human performance still exceeds the performance of state-of-the-art neural networks on distorted images, even if the display time is very short (100ms). Fine-tuning with distorted images reduces this gap, but nevertheless a gap still exists. The 100ms display time does not allow for eye movements or other higher level visual processes. This tells us that the human visual system is more efficient at processing “gist” level information than DNNs.

Note that our dataset is noticeably easier than the dataset in [7]. This was to ensure that the human subjects could still perform recognition at 100ms display times. A side-effect of this choice of dataset is that it is much easier for the deep neural networks to perform recognition. For example, for such simple classes it may be sufficient to use simple color information to recognize the images. A more difficult dataset may expose a larger difference between human and DNN performance.

These experiments provide a look into the bias and shortcomings of DNNs by comparing them with the human visual system. Future work could analyze other scenarios where humans still outperform DNNs. These studies can help motivate future work into more robust deep neural networks.

#### Acknowledgment

The authors would like to thank NVIDIA Corporation for the donation of a TITAN X GPU used in these experi-

ments.

#### References

- [1] T. J. Andrews and D. M. Coppola. Idiosyncratic characteristics of saccadic eye movements when viewing different visual environments. *Vision research*, 39(17):2947–2953, 1999. 1
- [2] T. Bachmann. Identification of spatially quantised tachistoscopic images of faces: How many pixels does it take to carry identity? *European Journal of Cognitive Psychology*, 3(1):87–103, 1991. 1
- [3] A. Borji and L. Itti. Human vs. computer in scene and object recognition. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 113–120, 2014. 1
- [4] C. F. Cadieu, H. Hong, D. L. Yamins, N. Pinto, D. Ardila, E. A. Solomon, N. J. Majaj, and J. J. DiCarlo. Deep neural networks rival the representation of primate IT cortex for core visual object recognition. *PLoS computational biology*, 10(12):e1003963, 2014. 2
- [5] Y. Chen, R. McBain, and D. Norton. Specific vulnerability of face perception to noise: A similar effect in schizophrenia patients and healthy individuals. *Psychiatry research*, 225(3):619–624, 2015. 1
- [6] S. Dodge and L. Karam. Understanding how image quality affects deep neural networks. *International Conference on Quality of Multimedia Experience (QoMEX)*, 2016. 1
- [7] S. Dodge and L. Karam. A study and comparison of human and deep learning recognition performance under visual distortions. *International Conference on Computer Communications and Networks*, 2017. 1, 2, 3, 4, 6
- [8] L. Fei-Fei, R. Fergus, and P. Perona. Learning generative visual models from few training examples: An incremental bayesian approach tested on 101 object categories. *Computer Vision and Image Understanding*, 2007. 2
- [9] F. Fleuret, T. Li, C. Dubout, E. K. Wampller, S. Yantis, and D. Geman. Comparing machines and humans on a visual categorization test. *Proceedings of the National Academy of Sciences*, 108(43):17621–17625, 2011. 1
- [10] R. Geirhos, D. H. J. Janssen, H. H. Schtt, J. R. M. Bethge, and F. A. Wichmann. Comparing deep neural networks against humans: object recognition when the signal gets weaker. *arXiv preprint arXiv:1706.06969*, 2017. 1, 2
- [11] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016. 4
- [12] C. Keysers, D.-K. Xiao, P. Földiák, and D. I. Perrett. The speed of sight. *Journal of cognitive neuroscience*, 13(1):90–101, 2001. 2
- [13] S. R. Kheradpisheh, M. Ghodrati, M. Ganjtabesh, and T. Masquelier. Deep networks can resemble human feed-forward vision in invariant object recognition. *Scientific reports*, 6:32672, 2016. 2
- [14] A. Oliva. Gist of the scene. *Neurobiology of attention*, 2005. 1

- [15] D. Parikh. Recognizing jumbled images: The role of local and global information in image classification. In *IEEE International Conference on Computer Vision (ICCV)*, pages 519–526, 2011. [1](#)
- [16] M. C. Potter, B. Wyble, C. E. Hagmann, and E. S. McCourt. Detecting meaning in rsvp at 13 ms per picture. *Attention, Perception, & Psychophysics*, 76(2):270–279, 2014. [2](#)
- [17] E. Rodner, M. Simon, R. B. Fisher, and J. Denzler. Fine-grained recognition in the noisy wild: Sensitivity analysis of convolutional neural networks approaches. *British Machine Vision Conference (BMVC)*, 2016. [1](#)
- [18] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015. [1](#), [2](#), [4](#)
- [19] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *International Conference on Learning Representations*, 2014. [4](#)
- [20] S. Stabinger, A. Rodríguez-Sánchez, and J. Piater. 25 years of cnns: Can we compare to human abstraction capabilities? In *International Conference on Artificial Neural Networks*, pages 380–387. Springer, 2016. [1](#)
- [21] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1–9, 2015. [4](#)
- [22] A. Torralba, R. Fergus, and W. T. Freeman. 80 million tiny images: A large data set for nonparametric object and scene recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 30(11):1958–1970, 2008. [1](#)
- [23] F. Wang, M. Jiang, C. Qian, S. Yang, C. Li, H. Zhang, X. Wang, and X. Tang. Residual attention network for image classification. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [24] A. R. Zamir, T.-L. Wu, L. Sun, W. Shen, J. Malik, and S. Savarese. Feedback networks. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.