

# Five points to check when comparing visual perception in humans and machines

**Christina M. Funke\***

University of Tübingen, Tübingen, Germany



**Judy Borowski\***

University of Tübingen, Tübingen, Germany



University of Tübingen, Tübingen, Germany  
Bernstein Center for Computational Neuroscience,  
Tübingen and Berlin, Germany  
Volkswagen Group Machine Learning Research Lab,  
Munich, Germany



**Karolina Stosio**

University of Tübingen, Tübingen, Germany  
Bernstein Center for Computational Neuroscience,  
Tübingen and Berlin, Germany  
Werner Reichardt Centre for Integrative Neuroscience,  
Tübingen, Germany



**Wieland Brendel†**

University of Tübingen, Tübingen, Germany  
Present address: Amazon.com, Tübingen



**Thomas S. A. Wallis†**

University of Tübingen, Tübingen, Germany  
Bernstein Center for Computational Neuroscience,  
Tübingen and Berlin, Germany

Werner Reichardt Centre for Integrative Neuroscience,  
Tübingen, Germany



**Matthias Bethge†**

With the rise of machines to human-level performance in complex recognition tasks, a growing amount of work is directed toward comparing information processing in humans and machines. These studies are an exciting chance to learn about one system by studying the other. Here, we propose ideas on how to design, conduct, and interpret experiments such that they adequately support the investigation of mechanisms when comparing human and machine perception. We demonstrate and apply these ideas through three case studies. The first case study shows how human bias can affect the interpretation of results and that several analytic tools can help to overcome this human reference point. In the second case study, we highlight the difference between necessary and sufficient mechanisms in visual reasoning tasks. Thereby, we show that contrary to previous suggestions, feedback mechanisms might not be necessary for the tasks in question. The third case study highlights the importance of aligning experimental conditions. We find that a previously observed

difference in object recognition does not hold when adapting the experiment to make conditions more equitable between humans and machines. In presenting a checklist for comparative studies of visual reasoning in humans and machines, we hope to highlight how to overcome potential pitfalls in design and inference.

## Introduction

Until recently, only biological systems could abstract the visual information in our world and transform it into a representation that supports understanding and action. Researchers have been studying how to implement such transformations in artificial systems since at least the 1950s. One advantage of artificial systems for understanding these computations is that many analyses can be performed that would not be possible in biological systems. For example, key

Citation: Funke, C. M., Borowski, J., Stosio, K., Brendel, W., Wallis, T. S. A., & Bethge, M. (2021). Five points to check when comparing visual perception in humans and machines. *Journal of Vision*, 21(3):16, 1–23, <https://doi.org/10.1167/jov.21.3.16>.



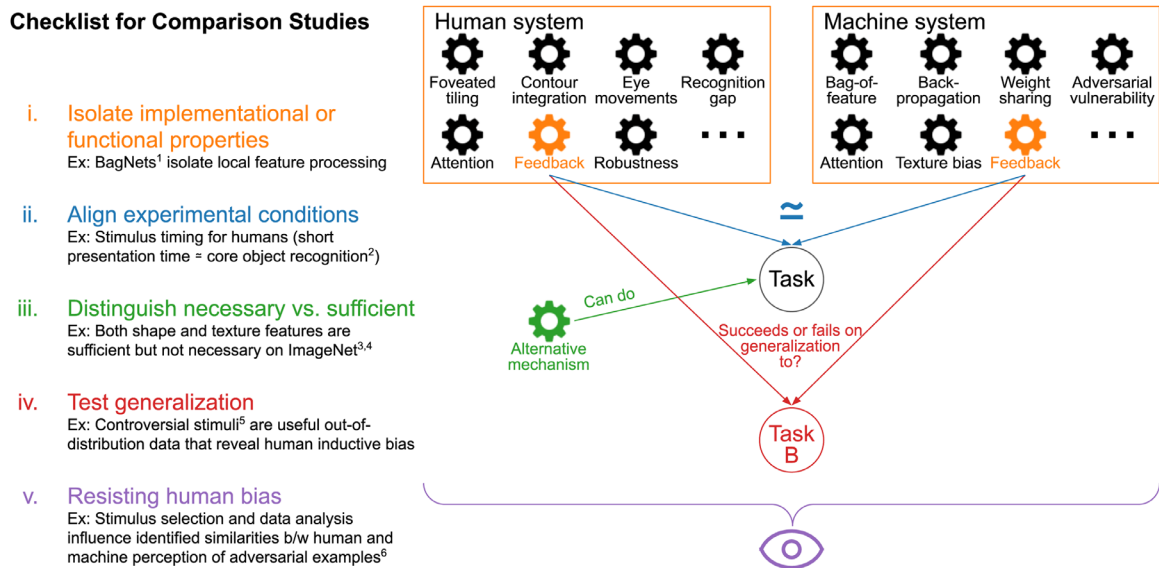


Figure 1. i: The human system and a candidate machine system differ in a range of properties. Isolating a specific mechanism (for example, feedback) can be challenging. ii: When designing an experiment, equivalent settings are important. iii: Even if a specific mechanism was important for a task, it would not be clear if this mechanism is necessary, as there could be other mechanisms (that might or might not be part of the human or machine system) that can allow a system to perform well. iv: Furthermore, the identified mechanisms might depend on the specific experimental setting and not generalize to, for example, another task. v: Overall, our human bias influences how we conduct and interpret our experiments. <sup>1</sup>Brendel and Bethge (2019); <sup>2</sup>DiCarlo et al. (2012); <sup>3</sup>Geirhos, Rubisch, et al. (2018); <sup>4</sup>Kubilius et al. (2016); <sup>5</sup>Golan et al. (2019); <sup>6</sup>Dujmović et al. (2020).

components of visual processing, such as the role of feedback connections, can be investigated, and methods such as ablation studies gain new precision.

Traditional models of visual processing sought to explicitly replicate the hypothesized computations performed in biological visual systems. One famous example is the hierarchical HMAX-model (Fukushima, 1980; Riesenhuber & Poggio, 1999). It instantiates mechanisms hypothesized to occur in primate visual systems, such as template matching and max operations, whose goal is to achieve invariance to position, scale, and translation. Crucially, though, these models never got close to human performance in real-world tasks.

With the success of learned approaches in the past decade, and particularly that of convolutional deep neural networks (DNNs), we now have much more powerful models. In fact, these models are able to perform a range of constrained image understanding tasks with human-like performance (Krizhevsky et al., 2012; Eigen & Fergus, 2015; Long et al., 2015).

While matching machine performance with that of the human visual system is a crucial step, the inner workings of the two systems can still be very different. We hence need to move beyond comparing accuracies to understand how the systems' mechanisms differ (Geirhos et al., 2020; Chollet, 2019; Ma & Peters, 2020; Firestone, 2020).

The range of frequently considered mechanisms is broad. They not only concern the architectural

level (such as feedback vs. feed-forward connections, lateral connections, foveated architectures or eye movements, ...), but also involve different learning schemes (back-propagation vs. spike-timing-dependent plasticity/Hebbian learning, ...) as well as the nature of the representations themselves (such as reliance on texture rather than shape, global vs. local processing, ...). For an overview of comparison studies, please see Appendix A.

## Checklist for psychophysical comparison studies

We present a checklist on how to design, conduct, and interpret experiments of comparison studies that investigate relevant mechanisms for visual perception. The diagram in Figure 1 illustrates the core ideas that we elaborate on below.

- Isolating implementational or functional properties.** Naturally, the systems that are being compared often differ in more than just one aspect, and hence pinpointing one single reason for an observed difference can be challenging. One approach is to design an artificial network constrained such that the mechanism of interest will show its effect as clearly as possible. An example of such an attempt

- is [Brendel and Bethge \(2019\)](#), which constrained models to process purely local information by reducing their receptive field sizes. Unfortunately, in many cases, it is almost impossible to exclude potential side effects from other experimental factors such as architecture or training procedure. Therefore, making explicit if, how, and where results depend on other experimental factors is important.
- ii. **Aligning experimental conditions for both systems.** In comparative studies (whether humans and machines, or different organisms in nature), it can be exceedingly challenging to make experimental conditions equivalent. When comparing the two systems, any differences should be made as explicit as possible and taken into account in the design and analysis of the study. For example, the human brain profits from lifelong experience, whereas a machine algorithm is usually limited to learning from specific stimuli of a particular task and setting. Another example is the stimulus timing used in psychophysical experiments, for which there is no direct equivalent in stateless algorithms. Comparisons of human and machine accuracies must therefore be considered with the temporal presentation characteristics of the experiment. These characteristics could be chosen based on, for example, a definition of the behavior of interest as that occurring within a certain time after stimulus onset (as for, e.g., “core object recognition”; [DiCarlo et al., 2012](#)). [Firestone \(2020\)](#) highlights that as aligning systems perfectly may not be possible due to different “hardware” constraints such as memory capacity, unequal performance of two systems might still arise despite similar competencies.
  - iii. **Differentiating between necessary and sufficient mechanisms.** It is possible that multiple mechanisms allow good task performance — for example, DNNs can use either shape or texture features to reach high performance on ImageNet ([Geirhos, Rubisch, et al., 2018](#); [Kubilius et al., 2016](#)). Thus, observing good performance for one mechanism does not imply that this mechanism is strictly necessary or that it is employed by the human visual system. As another example, [Watanabe et al. \(2018\)](#) investigated whether the rotating snakes illusion ([Kitaoka & Ashida, 2003](#); [Conway et al., 2005](#)) could be replicated in artificial neural networks. While they found that this was indeed the case, we argue that the mechanisms must be different from the ones used by humans, as the illusion requires small eye movements or blinks ([Hisakata & Murakami, 2008](#); [Kuriki et al., 2008](#)), while the artificial model does not emulate such biological processes.
  - iv. **Testing generalization of mechanisms.** Having identified an important mechanism, one needs to make explicit for which particular conditions (class of tasks, data sets, ...) the conclusion is intended to hold. A mechanism that is important for one setup may or may not be important for another one. In other words, whether a mechanism works under generalized settings has to be explicitly tested. An example of outstanding generalization for humans is their visual *robustness* against various variations in the input. In DNNs, a mechanism to improve robustness is to “stylize” ([Gatys et al., 2016](#)) training data. First presented as raising performance on parametrically distorted images ([Geirhos, Rubisch, et al., 2018](#)), this mechanism was later shown to also improve performance on images suffering from common corruptions ([Michaelis et al., 2019](#)) but would be unlikely to help with adversarial robustness. From a different perspective, the work of [Golan et al. \(2019\)](#) on controversial stimuli is an example where using stimuli outside of the training distribution can be insightful. Controversial stimuli are synthetic images that are designed to trigger distinct responses for two machine models. In their experimental setup, the use of these out-of-distribution data allows the authors to reveal whether the inductive bias of humans is similar to one of the candidate models.
  - v. **Resisting human bias.** Human bias can affect not only the design but also the conclusions we draw from comparison experiments. In other words, our human reference point can influence, for example, how we interpret the behavior of other systems, be they biological or artificial. An example is the well-known Braitenberg vehicles ([Braitenberg, 1986](#)), which are defined by very simple rules. To a human observer, however, the vehicles’ behavior appears as arising from complex internal states such as fear, aggression, or love. This phenomenon of anthropomorphizing is well known in the field of comparative psychology ([Romanes, 1883](#); [Köhler, 1925](#); [Koehler, 1943](#); [Haun et al., 2010](#); [Boesch, 2007](#); [Tomasello & Call, 2008](#)). [Buckner \(2019\)](#) specifically warns of human-centered interpretations and recommends to apply the lessons learned in comparative psychology to comparing DNNs and humans. In addition, our human reference point can influence how we design an experiment. As an example, [Dujmović et al. \(2020\)](#) illustrate that the selection of stimuli and labels can have a big effect on finding similarities or differences between humans and machines to adversarial examples.
- In the remainder of this article, we provide concrete examples of the aspects discussed above using three case studies<sup>1</sup>:
- (1) **Closed contour detection:** The first case study illustrates how tricky overcoming our human bias

can be and that shedding light on an alternative decision-making mechanism may require multiple additional experiments.

- (2) **Synthetic Visual Reasoning Test:** The second case study highlights the challenge of isolating mechanisms and of differentiating between necessary and sufficient mechanisms. Thereby, we discuss how human and machine model learning differ and how changes in the model architecture can affect the performance.
- (3) **Recognition gap:** The third case study illustrates the importance of aligning experimental conditions.

## Case study 1: Closed contour detection

Closed contours play a special role in human visual perception. According to the Gestalt principles of prägnanz and good continuation, humans can group distinct visual elements together so that they appear as a “form” or “whole.” As such, closed contours are thought to be prioritized by the human visual system and to be important in perceptual organization (Koffka, 2013; Elder & Zucker, 1993; Kovacs & Julesz, 1993; Tversky et al., 2004; Ringach & Shapley, 1996). Specifically, to tell if a line closes up to form a closed contour, humans are believed to implement a process called “contour integration” that relies at least partially on global information (Levi et al., 2007; Loffler et al., 2003; Mathes & Fahle, 2007). Even many flanking, open contours would hardly influence humans’ robust closed contour detection abilities.

### Our experiments

We hypothesize that, in contrast to humans, closed contour detection is difficult for DNNs. The reason is that this task would presumably require long-range contour integration, but DNNs are believed to process mainly local information (Geirhos, Rubisch, et al., 2018; Brendel & Bethge, 2019). Here, we test how well humans and neural networks can separate closed from open contours. To this end, we create a custom data set, test humans and DNNs on it, and investigate the decision-making process of the DNNs.

### DNNs and humans reach high performance

We created a data set with two classes of images: The first class contained a closed contour; the second one did not. In order to make sure that the statistical properties of the two classes were similar, we included a main contour for both classes. While this contour

line closed up for the first class, it remained open for the second class. This main contour consisted of 3–9 straight-line segments. In order to make the task more difficult, we added several flankers with either one or two line segments that each had a length of at least 32 pixels (Figure 2A). The size of the images was  $256 \times 256$  pixels. All lines were black and the background was uniformly gray. Details on the stimulus generation can be found in Appendix B.

Humans identified the closed contour stimulus very reliably in a two-interval forced-choice task. Their performance was 88.39% ( $SEM = 2.96\%$ ) on stimuli whose generation procedure was identical to the training set. For stimuli with white instead of black lines, human participants reached a performance of 90.52% ( $SEM = 1.58\%$ ). The psychophysical experiment is described in Appendix B.

We fine-tuned a ResNet-50 (He et al., 2016) pretrained on ImageNet (Deng et al., 2009) on the closed contour data set. Similar to humans, it performed very well and reached an accuracy of 99.95% (see Figure 2A [i.i.d. to training]).

We found that both humans and our DNN reach high accuracy on the closed contour detection task. From a human-centered perspective, it is enticing to infer that the model had learned the concept of open and closed contours and possibly that it performs a similar contour integration-like process as humans. However, this would have been overhasty. To better understand the degree of similarity, we investigated how our model performs on variations of the data sets that were not used during the training procedure.

### Generalization tests reveal differences

Humans are expected to have no difficulties if the number of flankers, the color, or the shape of lines would differ. We here test our model’s robustness on such variants of the data set. If our model used similar decision-making processes as humans, it should be able to generalize well without any further training on the new images. This procedure is another perspective to shed light on whether our model really understood the concept of closedness or just picked up some statistical cues in the training data set.

We tested our model on 15 variants of the data set (out of distribution test sets) without fine-tuning on these variations. As shown in Figure 2A, B, our trained model generalized well to many but not all modified stimulus sets.

On the following variations, our model achieved high accuracy: Curvy contours (1, 3) were easily distinguishable for our model, as long as the diameter remained below 100 pixels. Also, adding a dashed, closed flanker (2) did not lower its performance. The classification ability of the model remained similarly high for the no-flankers (4) and the asymmetric



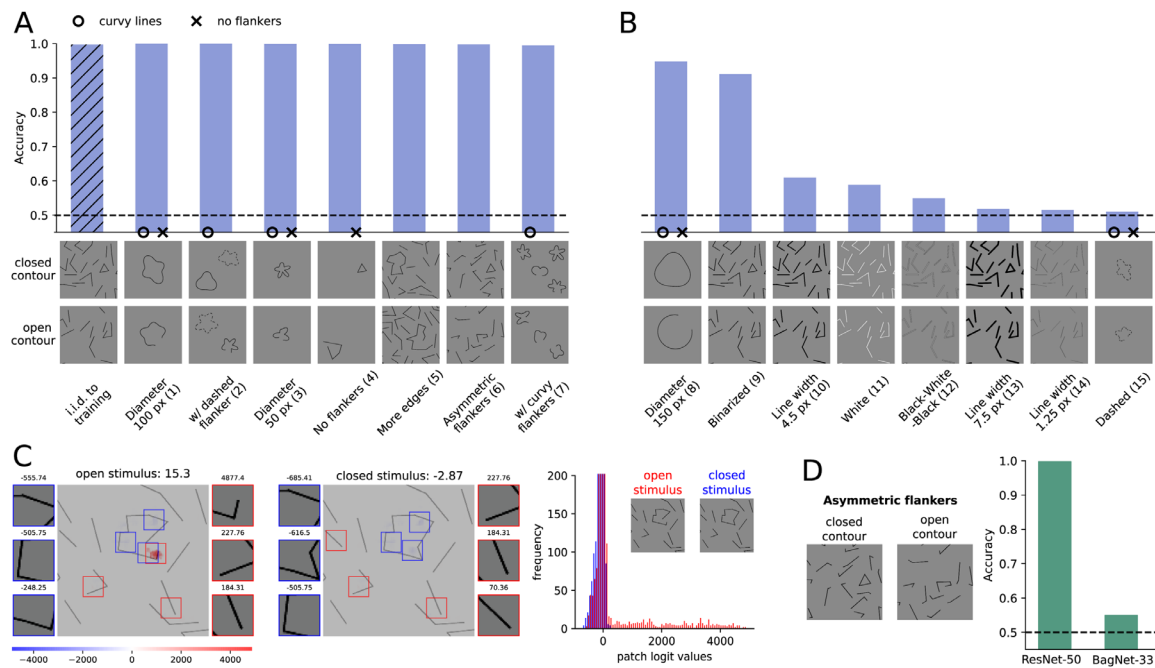


Figure 2. (A) Our ResNet-50-model generalized well to many data sets without further retraining, suggesting it would be able to distinguish closed and open contours. (B) However, the poor performance on many other data sets showed that our model did *not* learn the concept of closedness. (C) The heatmaps of our BagNet-33-based model show which parts of the image provided evidence for closedness (blue, negative values) or openness (red, positive values). The patches on the sides show the most extremely, nonoverlapping patches and their logit values. The logit distribution shows that most patches had logit values close to zero (y-axis truncated) and that many more patches in the open stimulus contributed positive logit values. (D) Our BagNet- and ResNet-models showed different performances on generalization sets, such as the asymmetric flankers. This indicates that the local decision-making process of the substitute model BagNet is not used by the original model ResNet. Figure best viewed electronically.

flankers condition (6). When testing our model on main contours that consisted of more edges than the ones presented during training (5), the performance was also hardly impaired. It remained high as well when multiple curvy open contours were added as flankers (7).

The following variations were more difficult for our model: If the size of the contour got too large, a moderate drop in accuracy was found (8). For binarized images, our model's performance was also reduced (9). And finally, (almost) chance performance was observed when varying the line width (14, 10, 13), changing the line color (11, 12), or using dashed curvy lines (15).

While humans would perform well on all variants of the closed contour data set, the failure of our model on some generalization tests suggests that it solves the task differently from humans. On the other hand, it is equally difficult to prove that the model does not understand the concept. As described by Firestone (2020), models can “perform differently despite similar underlying competences.” In either way, we argue that it is important to openly consider alternative mechanisms to the human approach of global contour integration.

## Our closed contour detection task is partly solvable with local features

In order to investigate an alternative mechanism to global contour integration, we here design an experiment to understand how well a decision-making process based on purely local features can work. For this purpose, we trained and tested BagNet-33 (Brendel & Bethge, 2019), a model that has access to local features only. It is a variation of ResNet-50 (He et al., 2016), where most  $3 \times 3$  kernels are replaced by  $1 \times 1$  kernels and therefore the receptive field size at the top-most convolutional layer is restricted to  $33 \times 33$  pixels.

We found that our restricted model still reached close to 90% performance. In other words, contour integration was not necessary to perform well on the task.

To understand which local features the model relied on mostly, we analyzed the contribution of each patch to the final classification decision. To this end, we used the log-likelihood values for each  $33 \times 33$  pixels patch from BagNet-33 and visualized them as a *heatmap*. Such a straightforward interpretation of the

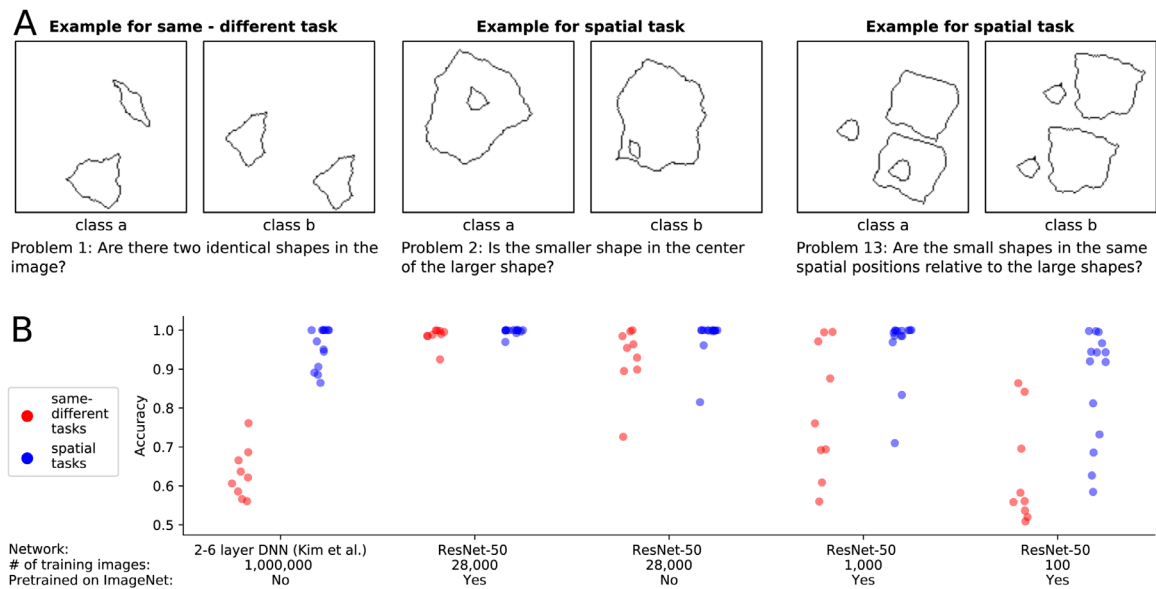


Figure 3. (A) For three of the 23 SVRT problems, two example images representing the two opposing classes are shown. In each problem, the task was to find the rule that separated the images and to sort them accordingly. (B) Kim et al. (2018) trained a DNN on each of the problems. They found that same-different tasks (red points), in contrast to spatial tasks (blue points), could not be solved with their models. Our ResNet-50-based models reached high accuracies for all problems when using 28,000 training examples and weights from pretraining on ImageNet.

contributions of single image patches is not possible with standard DNNs like ResNet (He et al., 2016) due to their large receptive field sizes in top layers.

The heatmaps of BagNet-33 (see Figure 2C) revealed which local patches played an important role in the decision-making process: An open contour was often detected by the presence of an endpoint at a short edge. Since all flankers in the training set had edges larger than 33 pixels, the presence of this feature was an indicator of an open contour. In turn, the absence of this feature was an indicator of a closed contour.

Whether the ResNet-50-based model used the same local feature as the substitute model was unclear. To answer this question, we tested BagNet on the previously mentioned generalization tests. We found that the data sets on which it showed high performance were sometimes different from the ones of ResNet (see Figure 7 in the Appendix B). A striking example was the failure of BagNet on the “asymmetric flankers” condition (see Figure 2D). For these images, the flankers often consisted of shorter line segments and thus obscured the local feature we assumed BagNet to use. In contrast, ResNet performed well on this variation. This suggests that the decision-making strategy of ResNet did not heavily depend on the local feature found with the substitute BagNet model.

In summary, the generalization tests, the high performance of BagNet as well as the existence of a distinctive local feature provide evidence that our human-biased assumption was misleading. We saw that

other mechanisms for closed contour detection besides global contour integration do exist (see Introduction, “Differentiating between necessary and sufficient mechanisms”). As humans, we can easily miss the many statistical subtleties by which a task can be solved. In this respect, BagNets proved to be a useful tool to test a purportedly “global” visual task for the presence of local artifacts. Overall, various experiments and analyses can be beneficial to understand mechanisms and to overcome our human reference point.

## Case study 2: Synthetic Visual Reasoning Test

In order to compare human and machine performance at learning abstract relationships between shapes, Fleuret et al. (2011) created the Synthetic Visual Reasoning Test (SVRT) consisting of 23 problems (see Figure 3A). They showed that humans need only few examples to understand the underlying concepts. Stabinger et al. (2016) as well as Kim et al. (2018) assessed the performance of deep convolutional neural networks on these problems. Both studies found a dichotomy between two task categories: While high accuracy was reached on spatial problems, the performance on same-different problems was poor. In order to compare the two types of tasks more systematically, Kim et al. (2018) developed a parameterized version of the SVRT data set called

PSVRT. Using this data set, they found that for same-different problems, an increase in the complexity of the data set could quickly strain their models. In addition, they showed that an attentive version of the model did not exhibit the same deficits. From these results, the authors concluded that feedback mechanisms as present in the human visual system such as attention, working memory, or perceptual grouping are probably important components for abstract visual reasoning. More generally, these studies have been perceived and cited with the broader claim of feed-forward DNNs not being able to learn same-different relationships between visual objects (Serre, 2019; Schofield et al., 2018) – at least not “efficiently” (Firestone, 2020).

We argue that the results of Kim et al. (2018) cannot be taken as evidence for the importance of feedback components for abstract visual reasoning:

- (1) While their experiments showed that same-different tasks are harder to *learn* for their models, this might also be true for the human visual system. Normally sighted humans have experienced lifelong visual input; only looking at human performance with this extensive learning experience cannot reveal differences in learning difficulty.
- (2) Even if there is a difference in learning complexity, this difference is not necessarily due to differences in the inference mechanism (e.g., feed-forward vs. feedback)—the large variety of other differences between biological and artificial vision systems could be critical causal factors as well.
- (3) In the same line, small modifications in the learning algorithm or architecture can significantly change learning complexity. For example, changing the network depth or width can greatly improve learning performance (Tan & Le, 2019).
- (4) Just because an attentive version of the model can learn both types of tasks does not prove that feedback mechanisms are necessary for these tasks (see Introduction, “*Differentiating between necessary and sufficient mechanisms*”).

Determining the necessity of feedback mechanisms is especially difficult because feedback mechanisms are not clearly distinct from purely feed-forward mechanisms. In fact, any finite-time recurrent network can be unrolled into a feed-forward network (Liao & Poggio, 2016; van Bergen & Kriegeskorte, 2020).

For these reasons, we argue that the importance of feedback mechanisms for abstract visual reasoning remains unclear.

In the following paragraph we present our own experiments on the SVRT data set and show that standard feed-forward DNNs can indeed perform well on same-different tasks. This confirms that feedback mechanisms are not strictly necessary for same-different tasks, although they helped in the specific experimental

setting of Kim et al. (2018). Furthermore, this experiment highlights that changes of the network architecture and training procedure can have large effects on the performance of artificial systems.

## Our experiments

The findings of Kim et al. (2018) were based on rather small neural networks, which consisted of up to six layers. However, typical network architectures used for object recognition consist of more layers and have larger receptive fields. For this reason, we tested a representative of such networks, namely, ResNet-50. The experimental setup can be found in Appendix C.

We found that our feed-forward model can in fact perform well on the same-different tasks of SVRT (see Figure 3B; see also concurrent work of Messina et al., 2019). This result was not due to an increase in the number of training samples. In fact, we used fewer images (28,000 images) than Kim et al. (2018) (1 million images) and Messina et al. (2019) (400,000 images). Of course, the results were obtained on the SVRT data set and might not hold for other visual reasoning data sets (see Introduction, “*Testing generalization of mechanisms*”).

In the very low-data regime (1,000 samples), we found a difference between the two types of tasks. In particular, the overall performance on same-different tasks was lower than on spatial reasoning tasks. As for the previously mentioned studies, this cannot be taken as evidence for systematic differences between feed-forward neural networks and the human visual system. In contrast to the neural networks used in this experiment, the human visual system is naturally pretrained on large amounts of visual reasoning tasks, thus making the low-data regime an unfair testing scenario from which it is almost impossible to draw solid conclusions about differences in the internal information processing. In other words, it might very well be that the human visual system trained from scratch on the two types of tasks would exhibit a similar difference in sample efficiency as a ResNet-50. Furthermore, the performance of a network in the low-data regime is heavily influenced by many factors other than architecture, including regularization schemes or the optimizer, making it even more difficult to reach conclusions about systematic differences in the network structure between humans and machines.

## Case study 3: Recognition gap

Ullman et al. (2016) investigated the minimally necessary visual information required for object recognition. To this end, they successively cropped or reduced the resolution of a natural image until more than 50% of all human participants failed to

identify the object. The study revealed that recognition performance drops sharply if the minimal recognizable image crops are reduced any further. They referred to this drop in performance as the “recognition gap.” The gap is computed by subtracting the proportion of people who correctly classify the largest unrecognizable crop (e.g., 0.2) from that of the people who correctly classify the smallest recognizable crop (e.g., 0.9). In this example, the recognition gap would evaluate to  $0.9 - 0.2 = 0.7$ . On the same human-selected image crops, Ullman et al. (2016) found that the recognition gap is much smaller for machine vision algorithms ( $0.14 \pm 0.24$ ) than for humans ( $0.71 \pm 0.05$ ). The researchers concluded that machine vision algorithms would not be able to “explain [humans] sensitivity to precise feature configurations” and “that the human visual system uses features and processes that are not used by current models and that are critical for recognition.” In a follow-up study, Srivastava et al. (2019) identified “fragile recognition images” (FRIs) with an exhaustive machine-based procedure whose results include a subset of patches that adhere to the definition of minimal recognizable configurations (MIRCs) by Ullman et al. (2016). On these machine-selected FRIs, a DNN experienced a moderately high recognition gap, whereas humans experienced a low one. Because of the differences between the selection procedures used in Ullman et al. (2016) and Srivastava et al. (2019), the question remained open whether machines would show a high recognition gap on machine-selected minimal images, if the selection procedure was similar to the one used in Ullman et al. (2016).

## Our experiment

Our goal was to investigate if the differences in recognition gaps identified by Ullman et al. (2016) would at least in part be explainable by differences in the experimental procedures for humans and machines. Crucially, we wanted to assess machine performance on *machine*-selected, and not *human*-selected, image crops. We therefore implemented the psychophysics experiment in a machine setting to search the smallest recognizable images (or MIRCs) and the largest unrecognizable images (sub-MIRCs). In the final step, we evaluated our machine model’s recognition gap using the *machine*-selected MIRCs and sub-MIRCs.

## Methods

Our machine-based search algorithm used the deep convolutional neural network BagNet-33 (Brendel & Bethge, 2019), which allows us to straightforwardly analyze images as small as  $33 \times 33$  pixels. In the first step, the classification accuracy was evaluated for the whole image. If it was above 0.5, the image was

successively cropped and reduced in resolution. In each step, the best-performing crop was taken as the new parent. When the classification probability of all children fell below 0.5, the parent was identified as the MIRC, and all its children were considered sub-MIRCs. In order to evaluate the recognition gap, we calculate the difference in accuracy between the MIRC and the *best-performing* sub-MIRC. This definition is more conservative than the one from Ullman et al. (2016), who evaluated the difference in accuracy between the MIRC and the *worst-performing* sub-MIRC. For more details on the search procedure, please see Appendix D.

## Results

We evaluated the recognition gap on two data sets: the original images from Ullman et al. (2016) and a subset of the ImageNet validation images (Deng et al., 2009). As shown in Figure 4A, our model has an average recognition gap of  $0.99 \pm 0.01$  on the machine-selected crops of the data set from Ullman et al. (2016). On the machine-selected crops of the ImageNet validation subset, a large recognition gap occurs as well. Our values are similar to the recognition gap in humans and differ from the machines’ recognition gap ( $0.14 \pm 0.24$ ) between human-selected MIRCs and sub-MIRCs as identified by Ullman et al. (2016).

## Discussion

Our findings contrast claims made by Ullman et al. (2016). The latter study concluded that machine algorithms are not as sensitive as humans to precise feature configurations and that they are missing features and processes that are “critical for recognition.” First, our study shows that a machine algorithm *is* sensitive to small image crops. It is only the precise minimal features that differ between humans and machines. Second, by the word “critical,” Ullman et al. (2016) imply that object recognition would not be possible without these human features and processes. Applying the same reasoning to Srivastava et al. (2019), the low human performance on machine-selected patches should suggest that humans would miss “features and processes critical for recognition.” This would be an obviously overreaching conclusion. Furthermore, the success of modern artificial object recognition speaks against the conclusion that the purported processes are “critical” for recognition, at least within this discretely defined recognition task. Finally, what we can conclude from the experiments of Ullman et al. (2016) and from our own is that both the human and a machine visual system can recognize small image crops and that there is a sudden drop in recognizability when reducing the amount of information.

In summary, these results highlight the importance of testing humans and machines in as similar settings



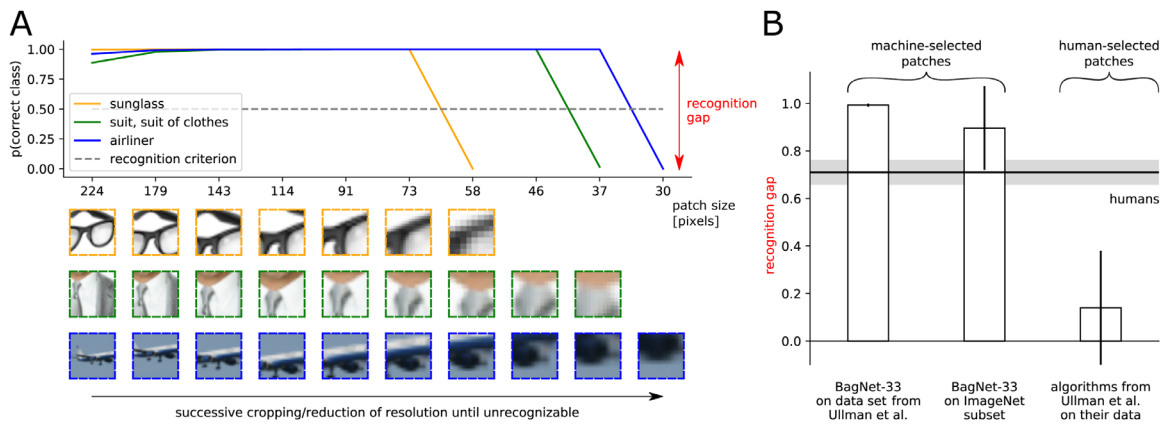


Figure 4. (A) BagNet-33’s probability of correct class for decreasing crops: The sharp drop when the image becomes too small or the resolution too low is called the “recognition gap” (Ullman et al., 2016). It was computed by subtracting the model’s predicted probability of the correct class for the sub-MIRC from the model’s predicted probability of the correct class for the MIRC. As an example, the glasses stimulus was evaluated as  $0.9999 - 0.0002 = 0.9997$ . The crop size on the x-axis corresponds to the size of the original image in pixels. Steps of reduced resolution are not displayed such that the three sample stimuli can be displayed coherently. (B) Recognition gaps for machine algorithms (vertical bars) and humans (gray horizontal bar). A recognition gap is identifiable for the DNN BagNet-33 when testing machine-selected stimuli of the original images from Ullman et al. (2016) and a subset of the ImageNet validation images (Deng et al., 2009). Error bars denote standard deviation.

as possible, and of avoiding a human bias in the experiment design. All conditions, instructions, and procedures should be as close as possible between humans and machines in order to ensure that observed differences are due to inherently different decision strategies rather than differences in the testing procedure.

## Conclusion

Comparing human and machine visual perception can be challenging. In this work, we presented a checklist on how to perform such comparison studies in a meaningful and robust way. For one, isolating a single mechanism requires us to minimize or exclude the effect of other differences between biological and artificial and to align experimental conditions for both systems. We further have to differentiate between necessary and sufficient mechanisms and to circumscribe in which tasks they are actually deployed. Finally, an overarching challenge in comparison studies between humans and machines is our strong internal human interpretation bias.

Using three case studies, we illustrated the application of the checklist. The first case study on closed contour detection showed that human bias can impede the objective interpretation of results and that investigating which mechanisms could or could not be at work may require several analytic tools. The second case study highlighted the difficulty of drawing robust conclusions about mechanisms from experiments. While previous studies suggested that feedback mechanisms might be

important for visual reasoning tasks, our experiments showed that they are not necessarily required. The third case study clarified that aligning experimental conditions for both systems is essential. When adapting the experimental settings, we found that, unlike the differences reported in a previous study, DNNs and humans indeed show similar behavior on an object recognition task.

Our checklist complements other recent proposals about how to compare visual inference strategies between humans and machines (Buckner, 2019; Chollet, 2019; Ma & Peters, 2020; Geirhos et al., 2020) and helps to create more nuanced and robust insights into both systems.

## Acknowledgments

The authors thank Alexander S. Ecker, Felix A. Wichmann, Matthias Kümmerer, Dylan Paiton, and Drew Linsley for helpful discussions. We thank Thomas Serre, Junkyung Kim, Matthew Ricci, Justus Piater, Sebastian Stabinger, Antonio Rodríguez-Sánchez, Shimon Ullman, Liav, Assif, and Daniel Harari for discussions and feedback on an earlier version of this manuscript. Additionally, we thank Nikolas Kriegeskorte for his detailed and constructive feedback, which helped us make our manuscript stronger. Furthermore, we thank Wiebke Ringels for helping with data collection for the psychophysical experiment.

We thank the International Max Planck Research School for Intelligent Systems (IMPRS-IS) for

supporting CMF and JB. We acknowledge support from the German Federal Ministry of Education and Research (BMBF) through the competence center for machine learning (FKZ 01IS18039A) and the Bernstein Computational Neuroscience Program Tübingen (FKZ: 01GQ1002), the German Excellence Initiative through the Centre for Integrative Neuroscience Tübingen (EXC307), and the Deutsche Forschungsgemeinschaft (DFG; Projektnummer 276693517 SFB 1233).

The closed contour case study was designed by CMF, JB, TSAW, and MB and later with WB. The code for the stimuli generation was developed by CMF. The neural networks were trained by CMF and JB. The psychophysical experiments were performed and analyzed by CMF, TSAW, and JB. The SVRT case study was conducted by CMF under supervision of TSAW, WB, and MB. KS designed and implemented the recognition gap case study under the supervision of WB and MB; JB extended and refined it under the supervision of WB and MB. The initial idea to unite the three projects was conceived by WB, MB, TSAW, and CMF, and further developed including JB. The first draft was jointly written by JB and CMF with input from TSAW and WB. All authors contributed to the final version and provided critical revisions.

Elements of this work were presented at the Conference on Cognitive Computational Neuroscience 2019 and the Shared Visual Representations in Human and Machine Intelligence Workshop at the Conference on Neural Information Processing Systems 2019.

The icon image is modified from the image by Gerd Leonhard, available under <https://www.flickr.com/photos/gleonhard/33661762360> on December 17, 2020. The original license is CC BY-SA 2.0, and therefore so is the one for the icon image.

Commercial relationships: Matthias Bethge: Amazon scholar Jan 2019 – Jan 2021, Layer7AI, DeepArt.io, Upload AI; Wieland Brendel: Layer7AI.

Corresponding authors: Christina M. Funke; Judy Borowski.

Email: [christina.funke@bethgelab.org](mailto:christina.funke@bethgelab.org); [judy.borowski@bethgelab.org](mailto:judy.borowski@bethgelab.org).

Address: Maria-von-Linden-Strasse 6, 72076, Tübingen, Germany.

\*CMF and JB are both first authors on this work.

†WB, TSAW and MB are joint senior authors.

## References

- Barrett, D. G., Hill, F., Santoro, A., Morcos, A. S., & Lillicrap, T. (2018). Measuring abstract reasoning in neural networks. In J. Dy, & A. Krause (Eds.), *Proceedings of the 35th International Conference on Machine Learning*, 80, 511–520. PMLR, <http://proceedings.mlr.press/v80/barrett18a.html>.
- Barrett, D. G., Morcos, A. S., & Macke, J. H. (2019). Analyzing biological and artificial neural networks: Challenges with opportunities for synergy? *Current Opinion in Neurobiology*, 55, 55–64.
- Boesch, C. (2007). What makes us human (homo sapiens)? The challenge of cognitive cross-species comparison. *Journal of Comparative Psychology*, 121(3), 227.
- Brainard, D. H., & Vision, S. (1997). The psychophysics toolbox. *Spatial Vision*, 10, 433–436.
- Braitenberg, V. (1986). *Vehicles: Experiments in synthetic psychology*. MIT Press, <https://mitpress.mit.edu/books/vehicles>.
- Brendel, W., & Bethge, M. (2019). Approximating CNNs with bag-of-local-features models works surprisingly well on imagenet. arXiv preprint, arXiv:1904.00760.
- Buckner, C. (2019). The comparative psychology of artificial intelligences, <http://philsci-archive.pitt.edu/16128/>.
- Cadena, S. A., Denfield, G. H., Walker, E. Y., Gatys, L. A., Tolias, A. S., Bethge, M., . . . Ecker, A. S. (2019). Deep convolutional models improve predictions of macaque v1 responses to natural images. *PLoS Computational Biology*, 15(4), e1006897.
- Chollet, F. (2019). The measure of intelligence. arXiv preprint, arXiv:1911.01547.
- Cichy, R. M., & Kaiser, D. (2019). Deep neural networks as scientific models. *Trends in Cognitive Sciences*, 23(4), 305–317, <https://doi.org/10.1016/j.tics.2019.01.009>.
- Conway, B. R., Kitaoka, A., Yazdanbakhsh, A., Pack, C. C., & Livingstone, M. S. (2005). Neural basis for a powerful static motion illusion. *Journal of Neuroscience*, 25(23), 5651–5656.
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., & Fei-Fei, L. (2009). ImageNet: A large-scale hierarchical image database. In *IEEE Conference on Computer Vision and Pattern Recognition*. pp. 248–255, <https://doi.org/10.1109/CVPR.2009.5206848>.
- DiCarlo, J. J., Zoccolan, D., & Rust, N. C. (2012). How does the brain solve visual object recognition? *Neuron*, 73(3), 415–434.
- Doerig, A., Bornet, A., Choung, O. H., & Herzog, M. H. (2019). Crowding reveals fundamental

## Footnote

<sup>1</sup>The code is available at [https://github.com/bethgelab/notorious\\_difficulty\\_of\\_comparing\\_human\\_and\\_machine\\_perception](https://github.com/bethgelab/notorious_difficulty_of_comparing_human_and_machine_perception).

- differences in local vs. global processing in humans and machines. *bioRxiv*, 744268.
- Dujmović, M., Malhotra, G., & Bowers, J. S. (2020). What do adversarial images tell us about human vision? *eLife*, *9*, e55978. Retrieved from <https://doi.org/10.7554/eLife.55978>.
- Eberhardt, S., Cader, J. G., & Serre, T. (2016). How deep is the feature analysis underlying rapid visual categorization? In D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, & R. Garnett (Eds.), *Advances in neural information processing systems*, *29*, 1100–1108. Curran Associates, Inc, <https://proceedings.neurips.cc/paper/2016/file/42e77b63637ab381e8be5f8318cc28a2-Paper.pdf>.
- Eigen, D., & Fergus, R. (2015). Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture. In *IEEE International Conference on Computer Vision* (pp. 2650–2658), doi:10.1109/ICCV.2015.304.
- Elder, J., & Zucker, S. (1993). The effect of contour closure on the rapid discrimination of two-dimensional shapes. *Vision Research*, *33*(7), 981–991.
- Elsayed, G., Shankar, S., Cheung, B., Papernot, N., Kurakin, A., Goodfellow, I., . . . Sohl-Dickstein, J. (2018). Adversarial examples that fool both computer vision and time-limited humans. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, & R. Garnett (Eds.), *Advances in neural information processing systems*, *31*, 3910–3920. Curran Associates, Inc.
- Firestone, C. (2020). Performance vs. competence in human-machine comparisons. *Proceedings of the National Academy of Sciences*. Retrieved from <https://www.pnas.org/content/early/2020/10/13/1905334117>.
- Fleuret, F., Li, T., Dubout, C., Wampller, E. K., Yantis, S., & Geman, D. (2011). Comparing machines and humans on a visual categorization test. *Proceedings of the National Academy of Sciences*, *108*(43), 17621–17625.
- Fukushima, K. (1980). Neocognitron: A self-organizing neural network model for a mechanism of pattern recognition unaffected by shift in position. *Biological Cybernetics*, *36*, 193–202.
- Gatys, L. A., Ecker, A. S., & Bethge, M. (2016). Image style transfer using convolutional neural networks. In *IEEE Conference on Computer Vision and Pattern Recognition* (pp. 2414–2423), doi:10.1109/CVPR.2016.265.
- Geirhos, R., Meding, K., & Wichmann, F. A. (2020). Beyond accuracy: Quantifying trial-by-trial behaviour of CNNs and humans by measuring error consistency. *arXiv preprint*, arXiv:2006.16736.
- Geirhos, R., Rubisch, P., Michaelis, C., Bethge, M., Wichmann, F. A., & Brendel, W. (2018). Imagenet-trained cnns are biased towards texture; increasing shape bias improves accuracy and robustness. *arXiv preprint*, arXiv:1811.12231.
- Geirhos, R., Temme, C. R. M., Rauber, J., Schütt, H. H., Bethge, M., & Wichmann, F. A. (2018). Generalisation in humans and deep neural networks. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, & R. Garnett (Eds.), *Advances in neural information processing systems*, *31*, 7538–7550. Curran Associates, Inc, <https://proceedings.neurips.cc/paper/2018/file/0937fb5864ed06ffb59ae5f9b5ed67a9-Paper.pdf>.
- Golan, T., Raju, P. C., & Kriegeskorte, N. (2019). Controversial stimuli: Pitting neural networks against each other as models of human recognition. *arXiv preprint*, arXiv:1911.09288.
- Gomez-Villa, A., Martin, A., Vazquez-Corral, J., & Bertalmio, M. (2019). Convolutional neural networks can be deceived by visual illusions. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (pp. 12301–12309), doi:10.1109/CVPR.2019.01259.
- Guo, C., Pleiss, G., Sun, Y., & Weinberger, K. Q. (2017). On calibration of modern neural networks. *Proceedings of the 34th International Conference on Machine Learning*, *70*, 1321–1330.
- Han, C., Yoon, W., Kwon, G., Nam, S., & Kim, D. (2019). Representation of white-and black-box adversarial examples in deep neural networks and humans: A functional magnetic resonance imaging study. *arXiv preprint*, arXiv:1905.02422.
- Haun, D. B. M., Jordan, F. M., Vallortigara, G., & Clayton, N. S. (2010). Origins of spatial, temporal, and numerical cognition: Insights from comparative psychology. *Trends in Cognitive Sciences*, *14*(12), 552–560, <https://doi.org/10.1016/j.tics.2010.09.006>, <http://www.sciencedirect.com/science/article/pii/S1364661310002135>.
- He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition* (pp. 770–778), doi:10.1109/CVPR.2016.90.
- Hisakata, R., & Murakami, I. (2008). The effects of eccentricity and retinal illuminance on the illusory motion seen in a stationary luminance gradient. *Vision Research*, *48*(19), 1940–1948.
- Khaligh-Razavi, S.-M., & Kriegeskorte, N. (2014). Deep supervised, but not unsupervised, models



- may explain its cortical representation. *PLoS Computational Biology*, 10(11), e1003915.
- Kim, B., Reif, E., Wattenberg, M., & Bengio, S. (2019). Do neural networks show gestalt phenomena? An exploration of the law of closure. arXiv preprint, arXiv:1903.01069.
- Kim, J., Ricci, M., & Serre, T. (2018). Not-so-clevr: Learning same–different relations strains feedforward neural networks. *Interface Focus*, 8(4), 20180011.
- Kingma, D. P., & Ba, J. (2014). Adam: A method for stochastic optimization. arXiv preprint, arXiv:1412.6980.
- Kitaoka, A., & Ashida, H. (2003). Phenomenal characteristics of the peripheral drift illusion. *Vision*, 15(4), 261–262.
- Kleiner, M., Brainard, D., Pelli, D., Ingling, A., Murray, R., & Broussard, C., (2007). What's new in psychtoolbox-3. *Perception*, 36(14), 1.
- Koehler, O. (1943). Counting experiments on a common raven and comparative experiments on humans. *Zeitschrift für Tierpsychologie*, 5(3), 575–712.
- Koffka, K. (2013). *Principles of Gestalt psychology*. New York: Routledge.
- Köhler, W. (1925). *The mentality of apes*. New York, NY: Kegan Paul, Trench, Trubner & Co.
- Kovacs, I., & Julesz, B. (1993). A closed curve is much more than an incomplete one: Effect of closure in figure-ground segmentation. *Proceedings of the National Academy of Sciences*, 90(16), 7495–7497.
- Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). Imagenet classification with deep convolutional neural networks. In F. Pereira, C. J. C. Burges, L. Bottou, & K. Q. Weinberger (Eds.), *Advances in Neural Information Processing Systems*, 25, 1097–1105. Curran Associates, Inc, <https://proceedings.neurips.cc/paper/2012/file/c399862d3b9d6b76c8436e924a68c45b-Paper.pdf>.
- Kubilius, J., Bracci, S., & de Beeck, H. P. O. (2016). Deep neural networks as a computational model for human shape sensitivity. *PLoS Computational Biology*, 12(4), e1004896.
- Kuriki, I., Ashida, H., Murakami, I., & Kitaoka, A. (2008). Functional brain imaging of the rotating snakes illusion by fMRI. *Journal of Vision*, 8(10), 16, <https://doi.org/10.1167/8.6.64>.
- Levi, D. M., Yu, C., Kuai, S.-G., & Rislove, E. (2007). Global contour processing in amblyopia. *Vision Research*, 47(4), 512–524.
- Liao, Q., & Poggio, T. (2016). Bridging the gaps between residual learning, recurrent neural networks and visual cortex. arXiv preprint, arXiv:1604.03640.
- Loffler, G., Wilson, H. R., & Wilkinson, F. (2003). Local and global contributions to shape discrimination. *Vision Research*, 43(5), 519–530.
- Long, J., Shelhamer, E., & Darrell, T. (2015). Fully convolutional networks for semantic segmentation. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (pp. 3431–3440), doi:10.1109/CVPR.2015.7298965.
- Luo, L., Xiong, Y., Liu, Y., & Sun, X. (2019). Adaptive gradient methods with dynamic bound of learning rate. arXiv preprint, arXiv:1902.09843.
- Ma, W. J., & Peters, B. (2020). A neural network walks into a lab: towards using deep nets as models for human behavior. arXiv preprint, arXiv:2005.02181.
- Majaj, N. J., & Pelli, D. G. (2018). Deep learning—using machine learning to study biological vision. *Journal of Vision*, 18(13), 2, <https://doi.org/10.1167/18.13.2>.
- Mathes, B., & Fahle, M. (2007). Closure facilitates contour integration. *Vision Research*, 47(6), 818–827, <https://doi.org/10.1167/18.13.2>.
- McCoy, R. T., Pavlick, E., & Linzen, T. (2019). Right for the wrong reasons: Diagnosing syntactic heuristics in natural language inference. arXiv preprint, arXiv:1902.01007.
- Messina, N., Amato, G., Carrara, F., Falchi, F., & Gennaro, C. (2019). Testing deep neural networks on the same-different task. In *International Conference on Content-Based Multimedia Indexing (CBMI)* (pp. 1–6). IEEE, doi:10.1109/CBMI.2019.8877412.
- Michaelis, C., Mitzkus, B., Geirhos, R., Rusak, E., Bringmann, O., Ecker, A. S., Bethge, M., . . . Brendel, W. (2019). Benchmarking robustness in object detection: Autonomous driving when winter is coming. arXiv preprint, arXiv:1907.07484.
- Miller, G. A. (1995). Wordnet: A lexical database for english. *Communications of the ACM*, 38(11), 39–41.
- Niven, T., & Kao, H.-Y. (2019). Probing neural network comprehension of natural language arguments. arXiv preprint, arXiv:1907.07355.
- Pelli, D. G., & Vision, S. (1997). The videotoolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, 10, 437–442.
- Peterson, J. C., Abbott, J. T., & Griffiths, T. L. (2016). Adapting deep network features to capture psychological representations. arXiv preprint, arXiv:1608.02164.
- Riesenhuber, M., & Poggio, T. (1999). Hierarchical models of object recognition in cortex. *Nature Neuroscience*, 2(11), 1019–1025.
- Ringach, D. L., & Shapley, R. (1996). Spatial and temporal properties of illusory contours and amodal boundary completion. *Vision Research*, 36(19), 3037–3050.



- Ritter, S., Barrett, D. G., Santoro, A., & Botvinick, M. M. (2017). Cognitive psychology for deep neural networks: A shape bias case study. *Proceedings of the 34th International Conference on Machine Learning*, 70, 2940–2949.
- Romanes, G. J. (1883). *Animal intelligence*. D. Appleton, [https://books.google.de/books?hl=en&lr=&id=Vx8aAAAAYAAJ&oi=fnd&pg=PA1&dq=animal+intelligence+1883&ots=IUOqpa2YRA&sig=JIVJfeIN7HlireTKzBd2tdv8IzM&redir\\_esc=y#v=onepage&q=animal%20intelligence%201883&f=false](https://books.google.de/books?hl=en&lr=&id=Vx8aAAAAYAAJ&oi=fnd&pg=PA1&dq=animal+intelligence+1883&ots=IUOqpa2YRA&sig=JIVJfeIN7HlireTKzBd2tdv8IzM&redir_esc=y#v=onepage&q=animal%20intelligence%201883&f=false).
- Santoro, A., Raposo, D., Barrett, D. G., Malinowski, M., Pascanu, R., Battaglia, P., ... Lillicrap, T. (2017). A simple neural network module for relational reasoning. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, ... R. Garnett (Eds.), *Advances in Neural Information Processing Systems*, 30, 4967–4976. Curran Associates, Inc, <https://proceedings.neurips.cc/paper/2017/file/e6acf4b0f69f6f6e60e9a815938aa1ff-Paper.pdf>.
- Schofield, A. J., Gilchrist, I. D., Bloj, M., Leonardis, A., & Bellotto, N. (2018). Understanding images in biological and computer vision. *Interface Focus*, 8, <https://doi.org/10.1098/rsfs.2018.0027>
- Schrimpf, M., Kumbilius, J., Hong, H., Majaj, N. J., Rajalingham, R., Issa, E. B., ... Schmidt, K. (2018). Brain-score: Which artificial neural network for object recognition is most brain-like? *BioRxiv*, 407007.
- Serre, T. (2019). Deep learning: The good, the bad, and the ugly. *Annual Review of Vision Science*, 5, 399–426.
- Spoerer, C. J., McClure, P., & Kriegeskorte, N. (2017). Recurrent convolutional neural networks: A better model of biological object recognition. *Frontiers in Psychology*, 8, 1551.
- Srivastava, S., Ben-Yosef, G., & Boix, X. (2019). Minimal images in deep neural networks: Fragile object recognition in natural images. *arXiv preprint*, arXiv:1902.03227.
- Stabinger, S., Rodríguez-Sánchez, A., & Piater, J. (2016). 25 years of CNNs: Can we compare to human abstraction capabilities? In *International Conference on Artificial Neural Networks*. (pp. 380–387). Cham: Springer.
- Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., ... Fergus, R. (2013). Intriguing properties of neural networks. *arXiv preprint*, arXiv:1312.6199.
- Tan, M., & Le, Q. V. (2019). Efficientnet: Rethinking model scaling for convolutional neural networks. *arXiv preprint*, arXiv:1905.11946.
- Tomasello, M., & Call, J. (2008). Assessing the validity of ape-human comparisons: A reply to boesch (2007). *Journal of Comparative Psychology*, 122(4), 449–452. American Psychological Association.
- Tversky, T., Geisler, W. S., & Perry, J. S. (2004). Contour grouping: Closure effects are explained by good continuation and proximity. *Vision Research*, 44(24), 2769–2777.
- Ullman, S., Assif, L., Fetaya, E., & Harari, D. (2016). Atoms of recognition in human and computer vision. *Proceedings of the National Academy of Sciences*, 113(10), 2744–2749.
- van Bergen, R. S., & Kriegeskorte, N. (2020). Going in circles is the way forward: The role of recurrence in visual inference. *arXiv preprint*, arXiv:2003.12128.
- Villalobos, K. M., Stih, V., Ahmadinejad, A., Dozier, J., Francl, A., Azevedo, F., Sasaki, T., ... Boix, X. (2020). Do deep neural networks for segmentation understand insideness? <https://cbmm.mit.edu/publications/do-neural-networks-segmentation-understand-insideness>.
- Volokitin, A., Roig, G., & Poggio, T. A. (2017). Do deep neural networks suffer from crowding? In: I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, ... R. Garnett (Eds.), *Advances in Neural Information Processing Systems*, 30, 5628–5638. Curran Associates, Inc, <https://proceedings.neurips.cc/paper/2017/file/c61f571dbd2fb949d3fe5ae1608dd48b-Paper.pdf>.
- Watanabe, E., Kitaoka, A., Sakamoto, K., Yasugi, M., & Tanaka, K. (2018). Illusory motion reproduced by deep neural networks trained for prediction. *Frontiers in Psychology*, 9, 345.
- Wu, X., Zhang, X., & Du, J. (2019). Challenge of spatial cognition for deep learning. *arXiv preprint*, arXiv:1908.04396.
- Yamins, D. L., Hong, H., Cadieu, C. F., Solomon, E. A., Seibert, D., & DiCarlo, J. J. (2014). Performance-optimized hierarchical models predict neural responses in higher visual cortex. *Proceedings of the National Academy of Sciences*, 111(23), 8619–8624.
- Yan, Z., & Zhou, X. S. (2017). How intelligent are convolutional neural networks? *arXiv preprint*, arXiv:1709.06126.
- Zhang, R., Wu, J., Zhang, C., Freeman, W. T., & Tenenbaum, J. B. (2016). A comparative evaluation of approximate probabilistic simulation and deep neural networks as accounts of human physical scene understanding. *arXiv preprint*, arXiv:1605.01138.
- Zhang, X., Watkins, Y., & Kenyon, G. T. (2018). Can deep learning learn the principle of closed

contour detection? In G. Bebis, R. Boyle, B. Parvin, D. Koracin, M. Turek, S. Ramalingam, K. Xu, S. Lin, B. Alsallakh, J. Yang, E. Cuervo, . . . J. Ventura (Eds.), *International Symposium on Visual Computing* (pp. 455–460). Cham: Springer, [https://doi.org/10.1007/978-3-030-03801-4\\_40](https://doi.org/10.1007/978-3-030-03801-4_40).

Zhou, Z., & Firestone, C. (2019). Humans can decipher adversarial images. *Nature Communications*, *10*(1), 1334.

## Appendix A: Literature overview of comparison studies

A growing body of work discusses comparisons of humans and machines on a higher level. [Majaj and Pelli \(2018\)](#) provide a broad overview how machine learning can help vision scientists to study biological vision, while [Barrett et al. \(2019\)](#) review methods on how to analyze representations of biological and artificial networks. From the perspective of cognitive science, [Cichy and Kaiser \(2019\)](#) stress that deep learning models *can* serve as scientific models that not only provide both helpful predictions and explanations but that can also be used for exploration. Furthermore, from the perspective of psychology and philosophy, [Buckner \(2019\)](#) emphasizes often-neglected caveats when comparing humans and DNNs such as human-centered interpretations and calls for discussions regarding how to properly align machine and human performance. [Chollet \(2019\)](#) proposes a general artificial intelligence benchmark and suggests to rather evaluate intelligence as “skill-acquisition efficiency” than to focus on skills at specific tasks.

In the following, we give a brief overview of studies that compare human and machine perception. In order to test if DNNs have similar cognitive abilities as humans, a number of studies test DNNs on abstract (visual) reasoning tasks ([Barrett et al., 2018](#); [Yan & Zhou, 2017](#); [Wu et al., 2019](#); [Santoro et al., 2017](#); [Villalobos et al., 2020](#)). Other comparison studies focus on whether human visual phenomena such as illusions ([Gomez-Villa et al., 2019](#); [Watanabe et al., 2018](#); [Kim et al., 2019](#)) or crowding ([Volokitin et al., 2017](#); [Doerig et al., 2019](#)) can be reproduced in computational models. In the attempt to probe intuition in machine models, DNNs are compared to intuitive physics engines, that is, probabilistic models that simulate physical events ([Zhang et al., 2016](#)).

Other works investigate whether DNNs are sensible models of human perceptual processing. To this end, their prediction or internal representations are compared to those of biological systems, for example, to human and/or monkey behavioral representations ([Peterson et al., 2016](#); [Schrimpf et al., 2018](#); [Yamins et al., 2014](#); [Eberhardt et al., 2016](#); [Golan et al., 2019](#)), human fMRI representations ([Han et al., 2019](#);

[Khaligh-Razavi & Kriegeskorte, 2014](#)) or monkey cell recordings ([Schrimpf et al., 2018](#); [Khaligh-Razavi & Kriegeskorte, 2014](#); [Yamins et al., 2014](#); [Cadena et al., 2019](#)).

A great number of studies focus on manipulating tasks and/or models. Researchers often use generalization tests on data dissimilar to the training set ([Zhang et al., 2018](#); [Wu et al., 2019](#)) to test whether machines understood the underlying concepts. In other studies, the degradation of object classification accuracy is measured with respect to image degradations ([Geirhos et al., 2018](#)) or with respect to the type of features that play an important role for human or machine decision-making ([Geirhos, Rubisch, et al., 2018](#); [Brendel & Bethge, 2019](#); [Kubilius et al., 2016](#); [Ullman et al., 2016](#); [Ritter et al., 2017](#)). A lot of effort is being put into investigating whether humans are vulnerable to small, adversarial perturbations in images ([Elsayed et al., 2018](#); [Zhou & Firestone, 2019](#); [Han et al., 2019](#); [Dujmović et al., 2020](#)), as DNNs are shown to be ([Szegedy et al., 2013](#)). Similarly, in the field of natural language processing, a trend is to manipulate the data set itself by, for example, negating statements to test whether a trained model gains an understanding of natural language or whether it only picks up on statistical regularities ([Niven & Kao, 2019](#); [McCoy et al., 2019](#)).

Further work takes inspiration from biology or uses human knowledge explicitly in order to improve DNNs. [Spoerer et al. \(2017\)](#) found that recurrent connections, which are abundant in biological systems, allow for higher object recognition performance, especially in challenging situations such as in the presence of occlusions—in contrast to pure feed-forward networks. Furthermore, several researchers suggest ([Zhang et al., 2018](#); [Kim et al., 2018](#)) or show ([Wu et al., 2019](#); [Barrett et al., 2018](#); [Santoro et al., 2017](#)) that designing networks’ architecture or features with human knowledge is key for machine algorithms to successfully solve abstract (reasoning) tasks.

## Appendix B: Closed contour detection

### Data set

Each image in the training set contained a main contour, multiple flankers, and a background image. The main contour and flankers were drawn into an image of size  $1,028 \times 1,028$  pixels. The main contour and flankers could be straight or curvy lines, for which the generation processes are respectively described in the next two subsections. The lines had a default thickness of 10 pixels. We then resized the image to  $256 \times 256$  pixels using anti-aliasing to transform the black and white pixels into smoother lines that had

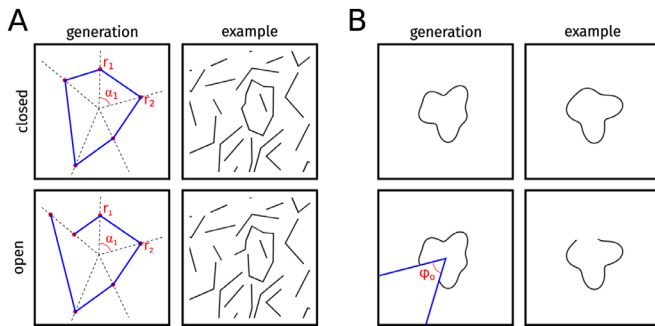


Figure 5. Closed contour data set. (A) Left: The main contour was generated by connecting points from a random sampling process of angles and radii. Right: Resulting line-drawing with flankers. (B) Left: Generation process of curvy contours. Right: Resulting line-drawing.

gray pixels at the borders. Thus, the lines in the resized image had a thickness of 2.5 pixels. In the following, all specifications of sizes refer to the resized image (i.e., a line described of final length 10 pixels extended over 40 pixels when drawn into the 1,028 × 1,028-pixel image). For the psychophysical experiments (see [Appendix B, Psychophysical experiment](#)), we added a white margin of 16 pixels on each side of the image to avoid illusory contours at the borders of the image.

**Varying contrast of background.** An image from the ImageNet data set was added as background to the line drawing. We converted the image into LAB color space and linearly rescaled the pixel intensities of the image to produce a normalized contrast value between 0 (gray image with the RGB values [118, 118, 118]) and 1 (original image) (see [Figure 8A](#)). When adding the image to the line drawing, we replaced all pixels of the line drawing by the values of the background image for which the background image had a higher grayscale value than the line drawing. For the experiments in the main body, the contrast of the background image was always 0. The other contrast levels were used only for the additional experiment described in [Appendix B, Additional experiment: Increasing the task difficulty by adding a background image](#).

**Generation of image pairs.** We aimed to reduce the statistical properties that could be exploited to solve the task without judging the closedness of the contour. Therefore, we generated image pairs consisting of an “open” and a “closed” version of the same image. The two versions were designed to be almost identical and had the same flankers. They differed only in the main contour, which was either open or closed. Examples of such image pairs are shown in [Figure 5](#). During training, either the closed or the open image of a pair was used. However, for the validation and testing, both versions were used. This allowed us to compare the predictions and heatmaps for images that differed only slightly but belonged to different classes.

### Line-drawing with polygons as main contour

The data set used for training as well as some of the generalization sets consisted of straight lines. The main contour consisted of  $n \in \{3, 4, 5, 6, 7, 8, 9\}$  line segments that formed either an open or a closed contour. The generation process of the main contour is depicted on the left side of [Figure 5A](#). To get a contour with  $n$  edges, we generated  $n$  points, which were defined by a randomly sampled angle  $\alpha_n$  and a randomly sampled radius  $r_n$  (between 0 and 128 pixels). By connecting the resulting points, we obtained the closed contour. We used the python PIL library (PIL 5.4.1, python3) to draw the lines that connect the endpoints. For the corresponding open contour, we sampled two radii for one of the angles such that they had a distance of 20 to 50 pixels from each other. When connecting the points, a gap was created between the points that share the same angle. This generation procedure could allow for very short lines with edges being very close to each other. To avoid this, we excluded all shapes with corner points closer to 10 pixels from nonadjacent lines.

The position of the main contour was random, but we ensured that the contour did not extend over the border of the image.

Besides the main contour, several flankers consisting of either one or two line segments were added to each stimulus. The exact number of flankers was uniformly sampled from the range [10,25]. The length of each line segment varied between 32 and 64 pixels. For the flankers consisting of two line segments, both lines had the same length, and the angle between the line segments was at least 45°. We added the flankers successively to the image and thereby ensured a minimal distance of 10 pixels between the line centers. To ensure that the corresponding image pairs would have the same flankers, the distances to both the closed and open version of the main contour were accounted for when re-sampling flankers. If a flanker did not fulfill this criterion, a new flanker was sampled of the same size and the same number of line segments, but it was placed somewhere else. If a flanker extended over the border of the image, the flanker was cropped.

### Line-drawing with curvy lines as main contour

For some of the generalization sets, the contours consisted of curvy instead of straight lines. These were generated by modulating a circle of a given radius  $r_c$  with a radial frequency function that was defined by two sinusoidal functions. The radius of the contour was thus given by

$$r(\phi) = A_1 \sin(f_1(\phi + \theta_1)) + A_2 \sin(f_2(\phi + \theta_2)) + r_c, \quad (1)$$

with the frequencies  $f_1$  and  $f_2$  (integers between 1 and 6), amplitudes  $A_1$  and  $A_2$  (random values between 15



and 45), and phases  $\theta_1$  and  $\theta_2$  (between 0 and  $2\pi$ ). Unless stated otherwise, the diameter (diameter =  $2 \times r_c$ ) was a random value between 50 and 100 pixels, and the contour was positioned in the center of the image. The open contours were obtained by removing a circular segment of size  $\phi_o = \frac{\pi}{3}$  at a random phase (see Figure 5B).

For two of the generalization data sets, we used dashed contours that were obtained by masking out 20 equally distributed circular segments each of size  $\phi_d = \frac{\pi}{20}$ .

### Details on generalization data sets

We constructed 15 variants of the data set to test generalization performance. Nine variants consisted of contours with straight lines. Six of these featured varying line styles like changes in line width (10, 13, 14) and/or line color (11, 12). For one variant (5), we increased the number of edges in the main contour. Another variant (4) had no flankers, and yet another variant (6) featured asymmetric flankers. For variant 9, the lines were binarized (only black or gray pixels instead of different gray tones).

In another six variants, the contours as well as the flankers were curved, meaning that we modulated a circle with a radial frequency function. The first four variants did not contain any flankers and the main contour had a fixed size of 50 pixels (3), 100 pixels (1), and 150 pixels (8). For another variant (15), the contour was a dashed line. Finally, we tested the effect of different flankers by adding one additional closed, yet dashed contour (2) or one to four open contours (7).

Below, we provide more details on some of these data sets:

**Black-white-black lines (12).** For all contours, black lines enclosed a white one in the middle. Each of these three lines had a thickness of 1.5 pixels, which resulted in a total thickness of 4.5 pixels.

**Asymmetric flankers (6).** The two-line flankers consisted of one long and one short line instead of two equally long lines.

**W/ dashed flanker (2).** This data set with curvy contours contained an additional dashed, yet closed contour as a flanker. It was produced like the main contour in the dashed main contour set. To avoid overlap of the contours, the main contour and the flanker could only appear at four determined positions in the image, namely, the corners.

**W/ multiple flankers (7).** In addition to the curvy main contour, between one and four open curvy contours were added as flankers. The flankers were generated by the same process as the main contour. The circles that were modulated had a diameter of 50 pixels and could appear at either one of the four corners of the image or in the center.

## Psychophysical experiment

To estimate how well humans would be able to distinguish closed and open stimuli, we performed a psychophysical experiment in which observers reported which of two sequentially presented images contained a closed contour (two-interval forced choice [2-IFC] task).

### Stimuli

The images of the closed contour data set were used as stimuli for the psychophysical experiments. Specifically, we used the images from the test sets that were used to evaluate the performance of the models. For our psychophysical experiments, we used two different conditions: The images contained either black (i.i.d. to the training set) or white contour lines. The latter was one of the generalization test sets.

### Apparatus

Stimuli were displayed on a VIEWPIXX 3D LCD (VIEWPIXX Technologies; spatial resolution 1,920 × 1,080 pixels, temporal resolution 120 Hz, operating with the scanning backlight turned off). Outside the stimulus image, the monitor was set to mean gray. Observers viewed the display from 60 cm (maintained via a chinrest) in a darkened chamber. At this distance, pixels subtended approximately  $0.024^\circ$  on average (41 pixels per degree of visual angle). The monitor was linearized (maximum luminance 260 cd/m<sup>2</sup> using a Konica-Minolta LS-100 photometer. Stimulus presentation and data collection were controlled via a desktop computer (Intel Core i5-4460 CPU, AMD Radeon R9 380 GPU) running Ubuntu Linux (16.04 LTS), using the Psychtoolbox Library (Pelli & Vision, 1997; Kleiner et al., 2007; Brainard & Vision, 1997, version 3.0.12) and the iShow library (<http://dx.doi.org/10.5281/zenodo.34217>) under MATLAB (The Mathworks, Inc., R2015b).

### Participants

In total, 19 naïve observers (4 male, 15 female, age: 25.05 years,  $SD = 3.52$ ) participated in the experiment. Observers were paid 10€ per hour for participation. Before the experiment, all subjects had given written informed consent for participating. All subjects had normal or corrected-to-normal vision. All procedures conformed to Standard 8 of the American Psychological Association's "Ethical Principles of Psychologists and Code of Conduct" (2010).

### Procedure

On each trial, one closed and one open contour stimulus were presented to the observer (see Figure 6A).



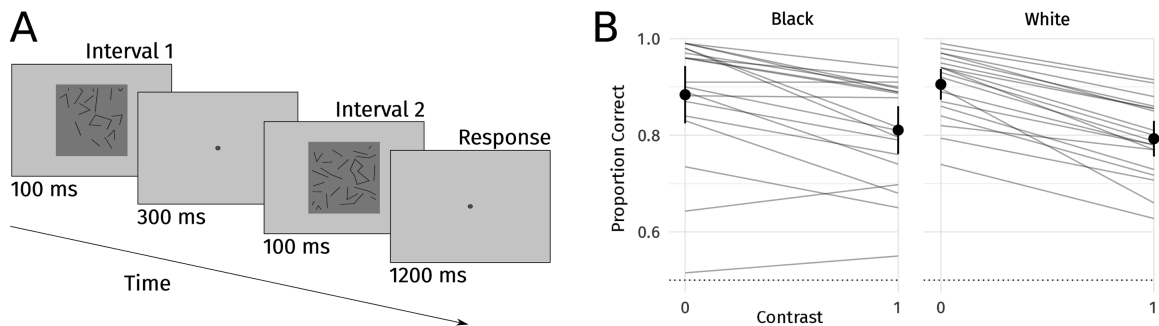


Figure 6. (A) In a 2-IFC task, human observers had to tell which of two images contained a closed contour. (B) Accuracy of the 20 naive observers for the different conditions.

The images used for each trial were randomly picked, but we ensured that the open and closed images shown in the same trial were not the ones that were almost identical to each other (see [Appendix B, Generation of image pairs](#)). Thus, the number of edges of the main contour could differ between the two images shown in the same trial. Each image was shown for 100 ms, separated by a 300-ms interstimulus interval (blank gray screen). We instructed the observer to look at the fixation spot in the center of the screen. The observer was asked to identify whether the image containing a closed contour appeared first or second. The observer had 1,200 ms to respond and was given feedback after each trial. The intertrial interval was 1,000 ms. Each block consisted of 100 trials and observers performed five blocks. Trials with different line colors and varying background images (contrasts including 0, 0.4, and 1) were blocked. Here, we only report the results for black and white lines of contrast 0. Upon the first time that a block with a new line color was shown, observers performed a practice session with 48 trials of the corresponding line color.

### Training of ResNet-50 model

We fine-tuned a ResNet-50 ([He et al., 2016](#)) pretrained on ImageNet ([Deng et al., 2009](#)) on the closed contour task. We replaced the last fully connected, 1,000-way classification layer by layer with only one output neuron to perform binary classification with a decision threshold of 0. The weights of all layers were fine-tuned using the optimizer Adam ([Kingma & Ba, 2014](#)) with a batch size of 64. All images were preprocessed to have the same mean and standard deviation and were randomly mirrored horizontally and vertically for data augmentation. The model was trained on 14,000 images for 10 epochs with a learning rate of 0.0003. We used a validation set of 5,600 images.

**Generalization tests.** To determine the generalization performance, we evaluated the model on the test sets

without any further training. Each of the test sets contained 5,600 images. Poor accuracy could simply result from a suboptimal decision criterion rather than because the network would not be able to tell the stimuli apart. To account for the distribution shift between the original training images and the generalization tasks, we optimized the decision threshold (a single scalar) for each data set. To find the optimal threshold for each data set, we subdivided the interval, in which 95% of all logits lie, into 100 sub points and picked the threshold that would lead to the highest performance.

### Training of BagNet-33 model

To test an alternative decision-making mechanism to global contour integration, we trained and tested a BagNet-33 ([Brendel & Bethge, 2019](#)) on the closed contour task. Like the ResNet-50 model, it was pretrained on ImageNet ([Deng et al., 2009](#)) and we replaced the last fully connected, 1,000-way classification layer by layer with only one output neuron. We fine-tuned the weights using the optimizer AdaBound ([Luo et al., 2019](#)) with an initial and final learning rate of 0.0001 and 0.1, respectively. The training images were generated on-the-fly, which meant that new images were produced for each epoch. In total, the fine-tuning lasted 100 epochs, and we picked the weights from the epoch with the highest performance.

**Generalization tests.** The generalization tests were conducted equivalently to the ones with ResNet-50. The results are shown in [Figure 7](#).

### Additional experiment: Increasing the task difficulty by adding a background image

We performed an additional experiment, where we tested if the model would become more robust and thus generalized better if we trained on a more difficult task. This was achieved by adding an image to the

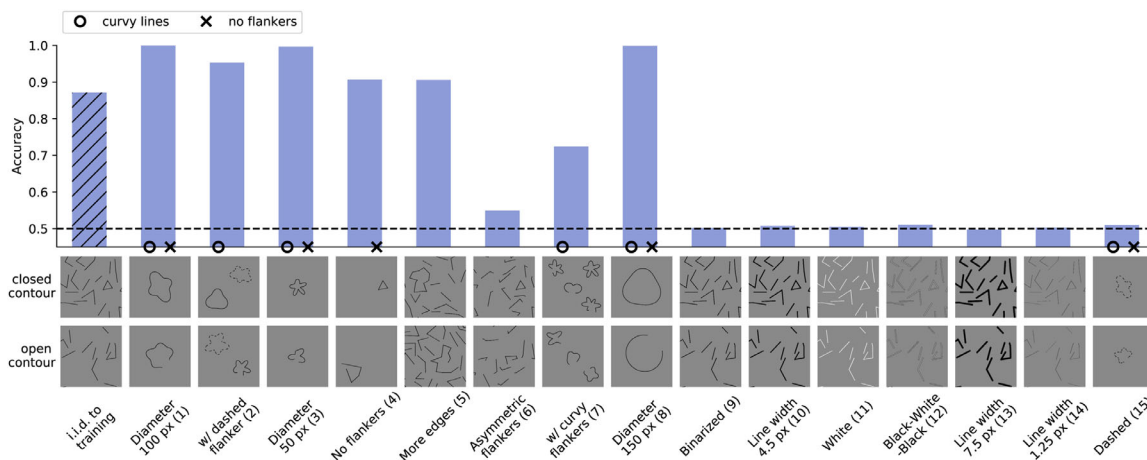


Figure 7. Generalization performances of BagNet-33.

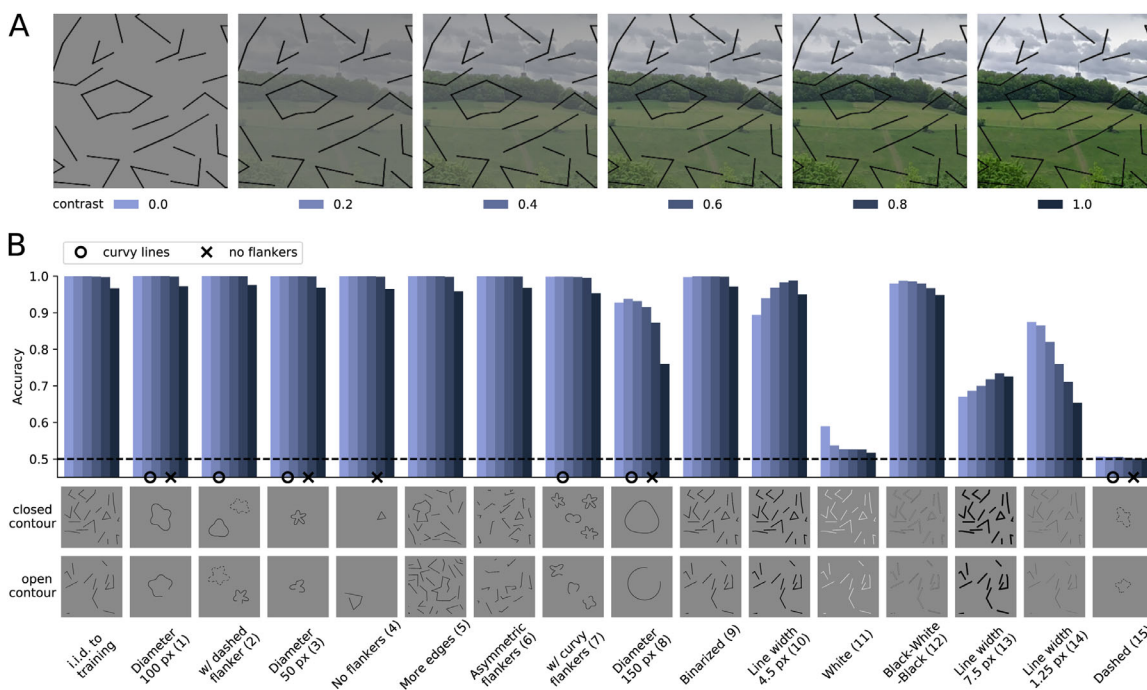


Figure 8. (A) An image of varying contrast was added as background. (B) Generalization performances of our models trained on random contrast levels and tested on single contrast levels.

background, such that the model had to learn how to separate the lines from the task-irrelevant background.

In our experiment, we fine-tuned our ResNet-50-based model on images with a background image of a uniformly sampled contrast. For each data set, we evaluated the model separately on six discrete contrast levels {0, 0.2, 0.4, 0.6, 0.8, 1} (see Figure 8A). We found that the generalization performance varied for some data sets compared to the experiment in the main body (see Figure 8B).

## Appendix C: SVRT

### Methods

**Data set.** We used the original C-code provided by Fleuret et al. (2011) to generate the images of the SVRT data set. The images had a size of 128 × 128 pixels. For each problem, we used up to 28,000 images for training,

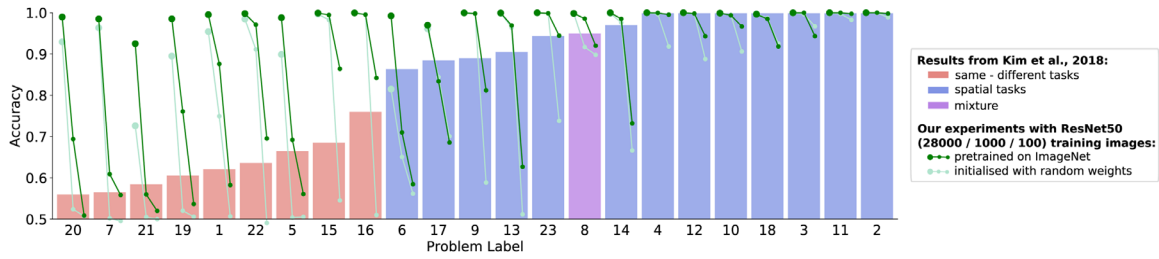


Figure 9. Accuracy of the models for the individual problems. Problem 8 is a mixture of same-different task and spatial task. In Figure 3, this problem was assigned to the spatial tasks. Bars replotted from Kim et al. (2018).

5,600 images for validation, and 11,200 images for testing.

**Experimental procedures.** For each of the SVRT problems, we fine-tuned a ResNet-50 that was pretrained on ImageNet (Deng et al., 2009) (as described in Appendix B, Training of ResNet-50 model). The same preprocessing, data augmentation, optimizer, and batch size as for the closed contour task were used.

For the different experiments, we varied the number of training images. We used subsets containing 28,000, 1,000, or 100 images. The number of epochs depended on the size of the training set: The model was fine-tuned for respectively 10, 280, or 2800 epochs. For each training set size and SVRT problem, we used the best learning rate after a hyper-parameter search on the validation set, where we tested the learning rates  $[6 \times 10^{-5}, 1 \times 10^{-4}, 3 \times 10^{-4}]$ .

As a control experiment, we also initialized the model with random weights, and we again performed a hyper-parameter search over the learning rates  $[3 \times 10^{-4}, 6 \times 10^{-4}, 1 \times 10^{-3}]$ .

## Results

In Figure 9, we show the results for the individual problems. When using 28,000 training images, we reached above 90% accuracy for all SVRT problems, including the ones that required same-different judgments (see also Figure 3B). When using less training images, the performance on the test set was reduced. In particular, we found that the performance on same-different tasks dropped more rapidly than on spatial reasoning tasks. If the ResNet-50 was trained from scratch (i.e., weights were randomly initialized instead of loaded from pretraining on ImageNet), the performance dropped only slightly on all but one spatial reasoning task. Larger drops were found on same-different tasks.

## Appendix D: Recognition gap

### Details on methods

**Data set.** We used two data sets for this experiment. One consisted of 10 natural, color images whose grayscale versions were also used in the original study by Ullman et al. (2016). We discarded one image from the original data set as it does not correspond to any ImageNet class. For our ground truth class selection, please see Table 1. The second data set consisted of 1,000 images from the ImageNet (Deng et al., 2009) validation set. All images were preprocessed like in standard training of ResNet (i.e., resizing to  $256 \times 256$  pixels, cropping centrally to  $224 \times 224$  pixels and normalizing).

**Model.** In order to evaluate the recognition gap, the model had to be able to handle small input images. Standard networks like ResNet (He et al., 2016) are not equipped to handle small images. In contrast, BagNet-33 (Brendel & Bethge, 2019) allows us to straightforwardly analyze images as small as  $33 \times 33$  pixels and hence was our model of choice for this experiment. It is a variation of ResNet-50 (He et al., 2016), where most  $3 \times 3$  kernels are replaced by  $1 \times 1$  kernels such that the receptive field size at the top-most convolutional layer is restricted to  $33 \times 33$  pixels.

**Machine-based search procedure for minimal recognizable images.** Similar to Ullman et al. (2016), we defined minimal recognizable images or configurations (MIRCs) as those patches of an image for which an observer—by which we mean an ensemble of humans or one or several machine algorithms—reaches  $\geq 50\%$  accuracy, but any additional 20% cropping of the corners or 20% reduction in resolution would lead to an accuracy  $< 50\%$ . MIRCs are thus inherently observer-dependent. The original study only searched for MIRCs in humans. We implemented the following procedure to find MIRCs in our DNN: We passed each preprocessed image through BagNet-33 and selected the most predictive crop according to its

Image	WordNet Hierarchy ID	WordNet Hierarchy description	Neuron number in ResNet-50 (indexing starts at 0)
fly	n02190166	fly	308
ship	n02687172	aircraft carrier, carrier, flattop, attack aircraft carrier	403
	n03095699	container ship, containership, container vessel	510
	n03344393	fireboat	554
	n03662601	lifeboat	625
	n03673027	liner, ocean liner	628
eagle	n01608432	kite	21
	n01614925	bald eagle, American eagle, <i>Haliaeetus leucocephalus</i>	22
	n04355933	sunglass	836
glasses	n04356056	sunglasses, dark glasses, shades	837
	n02835271	bicycle-built-for-two, tandem bicycle, tandem	444
bike	n03599486	jinrikisha, ricksha, rickshaw	612
	n03785016	moped	665
	n03792782	mountain bike, all-terrain bike, off-roader	671
	n04482393	tricycle, trike, velocipede	870
	n04350905	suit, suit of clothes	834
suit	n04591157	windsor tie	906
	n02690373	airliner	404
plane	n02389026	sorrel	339
horse	n03538406	horse cart, horse-cart	603
	n02701002	ambulance	407
car	n02814533	beach wagon, station wagon, wagon estate car, beach waggon, station waggon, waggon	436
	n02930766	cab, hack, taxi, taxicab	468
	n03100240	convertible	511
	n03594945	jeep, landrover	609
	n03670208	limousine, limo	627
	n03769881	minibus	654
	n03770679	minivan	656
	n04037443	racer, race car, racing car	751
	n04285008	sports car, sport car	817

Table 1. Selection of ImageNet classes for stimuli of Ullman et al. (2016).

probability. See [Appendix D, Selecting best crop when probabilities saturate](#) on how to handle cases where the probability saturates at 100% and [Appendix D, Analysis of different class selections and different number of descendants](#) for different treatments of ground truth class selections. If this probability of the full-size image for the ground-truth class was  $\geq 50\%$ , we again searched for the 80% subpatch with the highest probability. We repeated the search procedure until the class probability for all subpatches fell below 50%. If the 80% subpatches would be smaller than  $33 \times 33$  pixels, which is BagNet-33's smallest natural patch size, the crop was increased to  $33 \times 33$  pixels

using bilinear sampling. We evaluated the recognition gap as the difference in accuracy between the MIRC and the *best-performing* sub-MIRC. This definition was more conservative than the one from Ullman et al. (2016), who considered the maximum difference between a MIRC and its sub-MIRCs, that is, the difference between the MIRC and the *worst-performing* sub-MIRC. Please note that one difference between our machine procedure and the psychophysics experiment by Ullman et al. (2016) remained: The former was greedy, whereas the latter corresponded to an exhaustive search under certain assumptions.



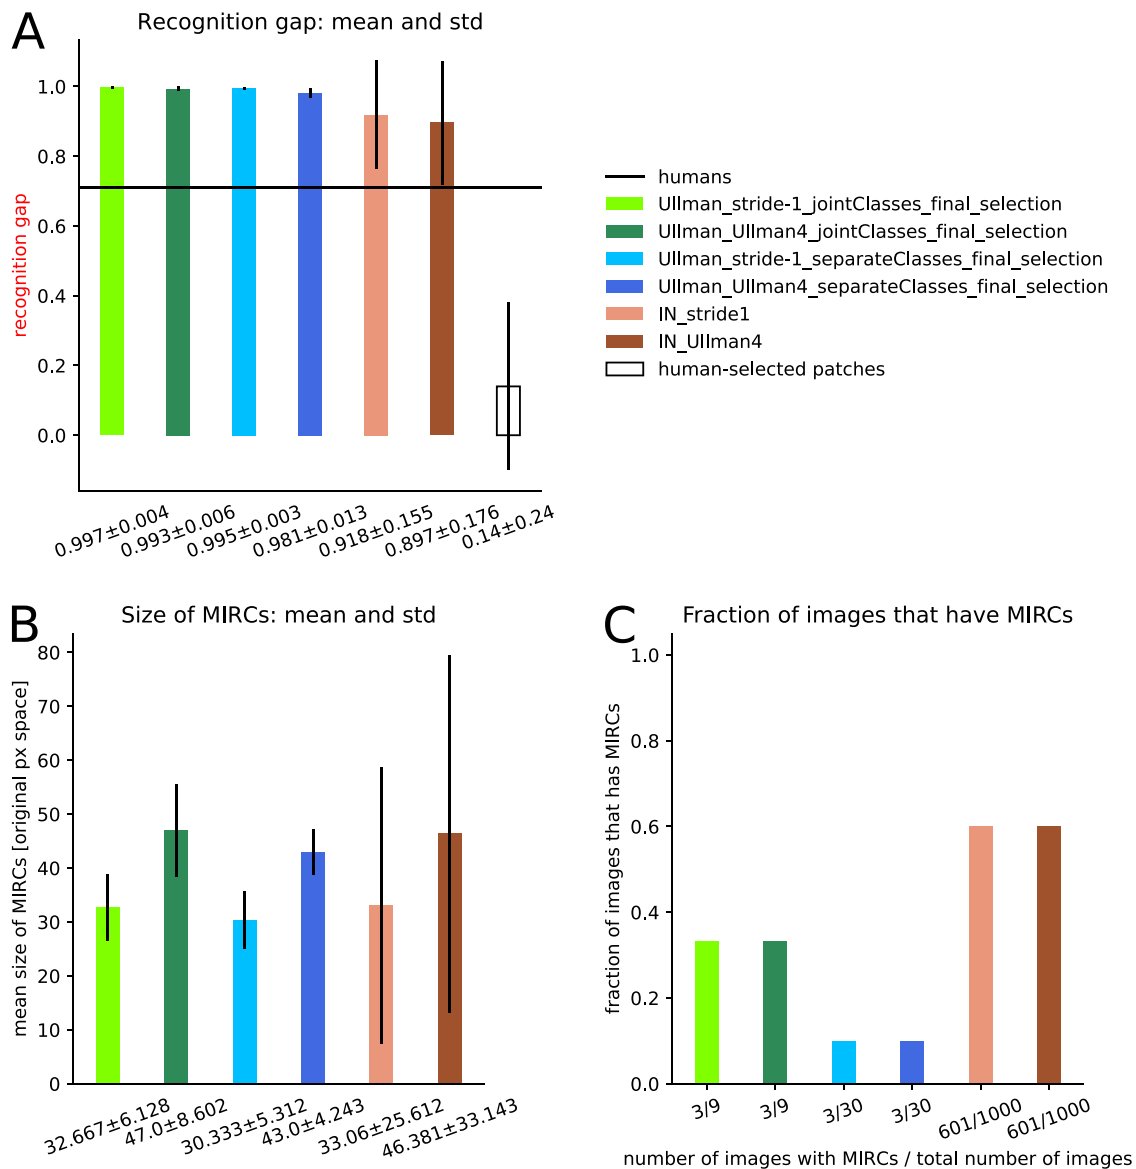


Figure 10. (A) Recognition gaps. The legend holds for all subplots. (B) Size of MIRCs. (C) Fraction of images that have MIRCs.

### Analysis of different class selections and different number of descendants

Treating the 10 stimuli from Ullman et al. (2016) in our machine algorithm setting required two design choices: We needed to both pick suitable ground truth classes from ImageNet for each stimulus as well as choose if and how to combine them. The former is subjective, and using relationships from WordNet Hierarchy (Miller, 1995) (as Ullman et al. [2016] did in their psychophysics experiment) only provides limited guidance. We picked classes to our best judgment (for our final ground truth class choices, please see Table 1). Regarding the aspect of handling several ground truth classes, we extended our experiments: We tested whether considering all classes as one (“joint classes,”

i.e., summing the probabilities) or separately (“separate classes,” i.e., rerunning the stimuli for each ground truth class) would have an effect on the recognition gap. As another check, we investigated whether the number of descendant options would alter the recognition gap: Instead of only considering the four corner crops as in the psychophysics experiment by Ullman et al. (2016) (“Ullman4”), we looked at every crop shifted by 1 pixel as a potential new parent (“stride-1”). The results reported in the main body correspond to joint classes and corner crops. Finally, besides analyzing the recognition gap, we also analyzed the sizes of MIRCs and the fractions of images that possess MIRCs for the mentioned conditions.

Figure 10A shows that all options result in similar values for the recognition gap. The trend of smaller

MIRC sizes for stride-1 compared to four corner crops shows that the search algorithm can find even smaller MIRCs when all crops are possible descendants (see [Figure 10B](#)). The final analysis of how many images possess MIRCs (see [Figure 10C](#)) shows that recognition gaps only exist for fractions of the tested images: In the case of the stimuli from [Ullman et al. \(2016\)](#), three out of nine images, and in the case of ImageNet, about 60% of the images have MIRCs. This means that the recognition performance of the initial full-size configurations was  $\geq 50\%$  for those fractions only. Please note that we did not evaluate the recognition gap over images that did not meet this criterion. In contrast, [Ullman et al. \(2016\)](#) average only across MIRCs that have a recognition rate above 65% and sub-MIRCs that have a recognition rate below 20% (personal communication, 2019). The reason why our model could only reliably classify three out of the nine stimuli from [Ullman et al. \(2016\)](#) can partly be traced back to the oversimplification of single-class attribution in ImageNet as well as to the overconfidence of deep learning classification algorithms ([Guo et al., 2017](#)): They often attribute a lot of evidence to one class, and the remaining ones only share very little evidence.

## Selecting best crop when probabilities saturate

We observed that several crops had very high probabilities and therefore used the “logit” measure  $\text{logit}(p)$ , where  $p$  is the probability. It is defined as the following:  $\text{logit}(p) = \log\left(\frac{p}{1-p}\right)$ . Note that this measure is different from what the deep learning community usually refers to as “logits,” which are the values before the softmax layer. In the following, we denote the latter values as  $\mathbf{z}$ . The logit  $\text{logit}(p)$  is monotonic w.r.t. to the probability  $p$ , meaning that the higher the probability  $p$ , the higher the logit  $\text{logit}(p)$ . However, while  $p$  saturates at 100%,  $\text{logit}(p)$  is unbounded. Therefore, it yields a more sensitive discrimination measure between image patches  $j$  that all have  $p(\mathbf{z}^j) = 1$ , where the superscript  $j$  denotes different patches.

In the following, we will provide a short derivation for the logit  $\text{logit}(p)$ . Consider a single patch with the correct class  $c$ . We start with the probability  $p_c$  of class  $c$ , which can be obtained by plugging the logits  $z_i$  into the softmax formula, where  $i$  corresponds to the classes  $[0, \dots, 1,000]$ .

$$p_c(\mathbf{z}) = \frac{\exp(z_c)}{\exp(z_c) + \sum_{i \neq c} \exp(z_i)} \quad (2)$$

Since we are interested in the probability of the correct class, it holds that  $p_c(\mathbf{z}) \neq 0$ . Thus, in the regime

of interest, we can invert both sides of the equation. After simplifying, we get

$$\frac{1}{p_c(\mathbf{z})} - 1 = \frac{\sum_{i \neq c} \exp(z_i)}{\exp(z_c)}. \quad (3)$$

When taking the negative logarithm on both sides, we obtain

$$\Leftrightarrow -\log\left(\frac{1}{p_c(\mathbf{z})} - 1\right) = -\log\left(\frac{\sum_{i \neq c} \exp(z_i)}{\exp(z_c)}\right) \quad (4)$$

$$\Leftrightarrow -\log\left(\frac{1 - p_c(\mathbf{z})}{p_c(\mathbf{z})}\right) = -\log\left(\sum_{i \neq c} \exp(z_i)\right) - (-\log(\exp(z_c))) \quad (5)$$

$$\Leftrightarrow \log\left(\frac{p_c(\mathbf{z})}{1 - p_c(\mathbf{z})}\right) = z_c - \log\left(\sum_{i \neq c} \exp(z_i)\right) \quad (6)$$

The left-hand side of the equation is exactly the definition of the logit  $\text{logit}(p)$ . Intuitively, it measures in log-space how much the network’s belief in the correct class outweighs the belief in all other classes taken together. The following reassembling operations illustrate this:

$$\begin{aligned} \text{logit}(p_c) &= \log\left(\frac{p_c(\mathbf{z})}{1 - p_c(\mathbf{z})}\right) \\ &= \underbrace{\log(p_c(\mathbf{z}))}_{\text{log probability of correct class}} \\ &\quad - \underbrace{\log(1 - p_c(\mathbf{z}))}_{\text{log probability of all incorrect classes}} \quad (7) \end{aligned}$$

The above formulations regarding one correct class hold when adjusting the experimental design to accept several classes  $k$  as correct predictions. In brief, the logit  $\text{logit}(p_C(\mathbf{z}))$ , where  $C$  stands for several classes, then states

$$\begin{aligned} \text{logit}(p_C(\mathbf{z})) &= -\log\left(\frac{1}{p_{c_1}(\mathbf{z}) + p_{c_2}(\mathbf{z}) + \dots + p_{c_k}(\mathbf{z})} - 1\right) \\ &= -\log\left(\frac{1}{\sum_k p_k(\mathbf{z})} - 1\right) \\ &= \underbrace{\log\left(\sum_k p_k(\mathbf{z})\right)}_{\text{log probability of all correct classes}} \end{aligned}$$

$$\begin{aligned}
 & - \underbrace{\log \left( 1 - \sum_k p_k(\mathbf{z}) \right)}_{\text{log probability of all incorrect classes}} \\
 & = \log \left( \sum_k \exp(z_k) \right) - \log \left( \sum_{i \neq k} \exp(z_i) \right) \quad (8)
 \end{aligned}$$

### Selection of ImageNet classes for stimuli of Ullman et al. (2016)

Note that our selection of classes is different from the one used by Ullman et al. (2016). We went through all classes for each image and selected the ones that we considered sensible. The 10th image of the eye does not have a sensible ImageNet class; hence, only nine stimuli from Ullman et al. (2016) are listed in Table 1.