

ARTICLE

<https://doi.org/10.1038/s41467-019-08931-6>

OPEN

Humans can decipher adversarial images

Zhenglong Zhou ¹ & Chaz Firestone¹

Does the human mind resemble the machine-learning systems that mirror its performance? Convolutional neural networks (CNNs) have achieved human-level benchmarks in classifying novel images. These advances support technologies such as autonomous vehicles and machine diagnosis; but beyond this, they serve as candidate models for human vision itself. However, unlike humans, CNNs are “fooled” by adversarial examples—nonsense patterns that machines recognize as familiar objects, or seemingly irrelevant image perturbations that nevertheless alter the machine’s classification. Such bizarre behaviors challenge the promise of these new advances; but do human and machine judgments fundamentally diverge? Here, we show that human and machine classification of adversarial images are robustly related: In 8 experiments on 5 prominent and diverse adversarial imagesets, human subjects correctly anticipated the machine’s preferred label over relevant foils—even for images described as “totally unrecognizable to human eyes”. Human intuition may be a surprisingly reliable guide to machine (mis)classification—with consequences for minds and machines alike.

¹ Department of Psychological & Brain Sciences, Johns Hopkins University, 3400 N Charles St., Baltimore, MD 21218, USA. Correspondence and requests for materials should be addressed to C.F. (email: chaz@jhu.edu)

How similar is the human mind to the machines that can behave like it? After decades spent lagging behind the recognitional capabilities of even a young child, machine-vision systems can now classify natural images with accuracy rates that match adult humans^{1,2}. The success of such models, especially biologically inspired Convolutional Neural Networks (CNNs³), has been exciting not only for the practical purpose of developing new technologies (e.g., screening baggage at airports, reading street signs in autonomous vehicles, or diagnosing radiological scans), but also for better understanding the human mind itself. Recent work, for example, has found that CNNs can be used to predict the behavior of humans and non-human primates, large-scale activation of brain regions, and even the firing patterns of individual neurons — leading to speculation that the mechanisms and computational principles underlying CNNs may resemble those of our own brains^{4–11}.

However, such models of object classification differ from humans in a crucial, alarming, and even bizarre way: They are vulnerable to attack by an “adversary”, such that a second model can be trained to produce images that “fool” the image-recognition model into misclassifying^{12–17}. Two especially striking classes of such adversarial images might be crudely called “fooling” images and “perturbed” images (Fig. 1). Fooling images are otherwise meaningless patterns that are classified as familiar objects by a machine-vision system. For example, a collection of oriented lines might be classified as a “baseball”, or a colorful television-static-like image might be called an “armadillo”. Perturbed images are images that would normally be classified accurately and straightforwardly (e.g., an ordinary photograph of a daisy, or a handwritten number 6) but that are perturbed only slightly to produce a completely different classification by the machine (e.g., a jaguar, or a handwritten number 5).

Adversarial images mark an “astonishing difference in the information processing of humans and machines”¹⁸, and fundamentally challenge the promise of these new approaches. First, and more practically, the existence of adversarial images could invite malicious attacks against machine-vision systems in applied settings (e.g., changing how an autonomous vehicle reads a street sign¹⁹). Second, and more theoretically, the fact that such bizarre images are straightforwardly classified as familiar objects by the machine seems to reveal how alien the CNN’s underlying processing must be^{11,20–22}, which in turn diminishes their utility as avenues for understanding the human mind.

A primary factor that makes adversarial images so intriguing is the intuitive assumption that a human would not classify the image as the machine does. (Indeed, this is part of what makes an image “adversarial” in the first place, though that definition is not yet fully settled.) However, surprisingly little work has actively explored this assumption by testing human performance on such images, even though it is often asserted that adversarial images are “totally unrecognizable to human eyes”¹⁴, p.427. At the same time, it has never been clear under which conditions human and machine performance might be usefully compared, especially since machine-vision systems of this sort typically have only a limited stock of labels to apply to such images. Indeed, you may informally observe that at least some adversarial images “make sense” once you are told the label assigned by the machine (as in Fig. 1). This raises an intriguing question: Could humans decipher such images by predicting the machine’s preferred labels? If so, this might suggest a greater overlap between human and machine classification than adversarial images seem to imply, and could even point to human intuition as a piece of the more practical puzzle of defending against such attacks.

To address this question, we introduce a “machine-theory-of-mind” task that asks whether humans can infer the classification that a machine-vision system would assign to a given image. We

acquired images produced by several prominent adversarial attacks, and displayed them to human subjects who were told that a machine had classified them as familiar objects. The human’s task was to “think like a machine” and determine which label was generated for each image. (For a related task with natural images, see ref. 23.) We conducted eight experiments using this task, probing human understanding of five different adversarial images^{12–15}. Importantly, none of these images was created with human vision in mind (cf.²⁴)—they were simply generated to fool a machine-vision system into misclassifying an image.

Across these 8 experiments—covering a diverse array of adversarial attacks, as well as several variations on the core experimental design—we find that human subjects can anticipate the machine’s classifications of adversarial stimuli. We conclude that human intuition is a more reliable guide to machine (mis) classification than has typically been imagined, and we discuss the implications of these results for comparisons between humans and machines in the context of cognitive science and artificial intelligence.

Results

Experiment 1: Fooling images with foil labels. Our first experiment administered the machine-theory-of-mind task using 48 “fooling” images that were produced by an evolutionary algorithm to confound a highly influential image-recognizing CNN, AlexNet¹, which classified them as familiar objects such as “pinwheel” and “bagel”. (There is evidence that adversarial images for one CNN often transfer to others²⁵.)

On each trial, subjects ($N=200$) saw one fooling image, displayed above both its CNN-generated label and a label randomly drawn from the other 47 images. Subjects selected whichever of the two labels they thought the machine generated for that image (Fig. 2a).

Remarkably, human observers strongly preferred the machine’s chosen labels to the foil labels: Classification “accuracy” (i.e., agreement with the machine’s classification) was 74%, well above-chance accuracy of 50% (95% confidence interval: [72.9, 75.5%]; two-sided binomial probability test: $p < 0.001$). Perhaps more tellingly, 98% of observers chose the machine’s label at above-chance rates, suggesting surprisingly universal agreement with the machine’s choices (Fig. 2d, “% of subjects who agree with the machine”). Additionally, 94% of the images showed above-chance human-machine agreement: Only 3 images out of all 48 had corresponding CNN-generated labels that humans tended to reject compared to a random label, whereas 45/48 had CNN-generated labels that the humans tended to prefer over a random label (Fig. 2d, “% of images with human-machine agreement”). This initial result suggests that human observers can broadly distinguish the features CNNs use to classify fooling images as familiar objects.

Experiment 2: 1st choice vs. 2nd choice. How deep does this ability run? Though subjects in Experiment 1 could distinguish the machine’s chosen label from a random label, they may have achieved this reliable classification not by discerning any meaningful resemblance between the images and their CNN-generated labels, but instead by identifying very superficial commonalities between them (e.g., preferring “bagel” to “pinwheel” for an orange-yellow blob simply because bagels are also orange-yellow in color).

To ask whether humans can appreciate subtler distinctions made by the machine, Experiment 2 contrasted the CNN’s preferred label not with random labels but instead with the machine’s second-ranked label for that image. For example, considering the round golden blob in Fig. 2, AlexNet’s next choice after “bagel” is “pretzel”, which similarly implies a curved golden

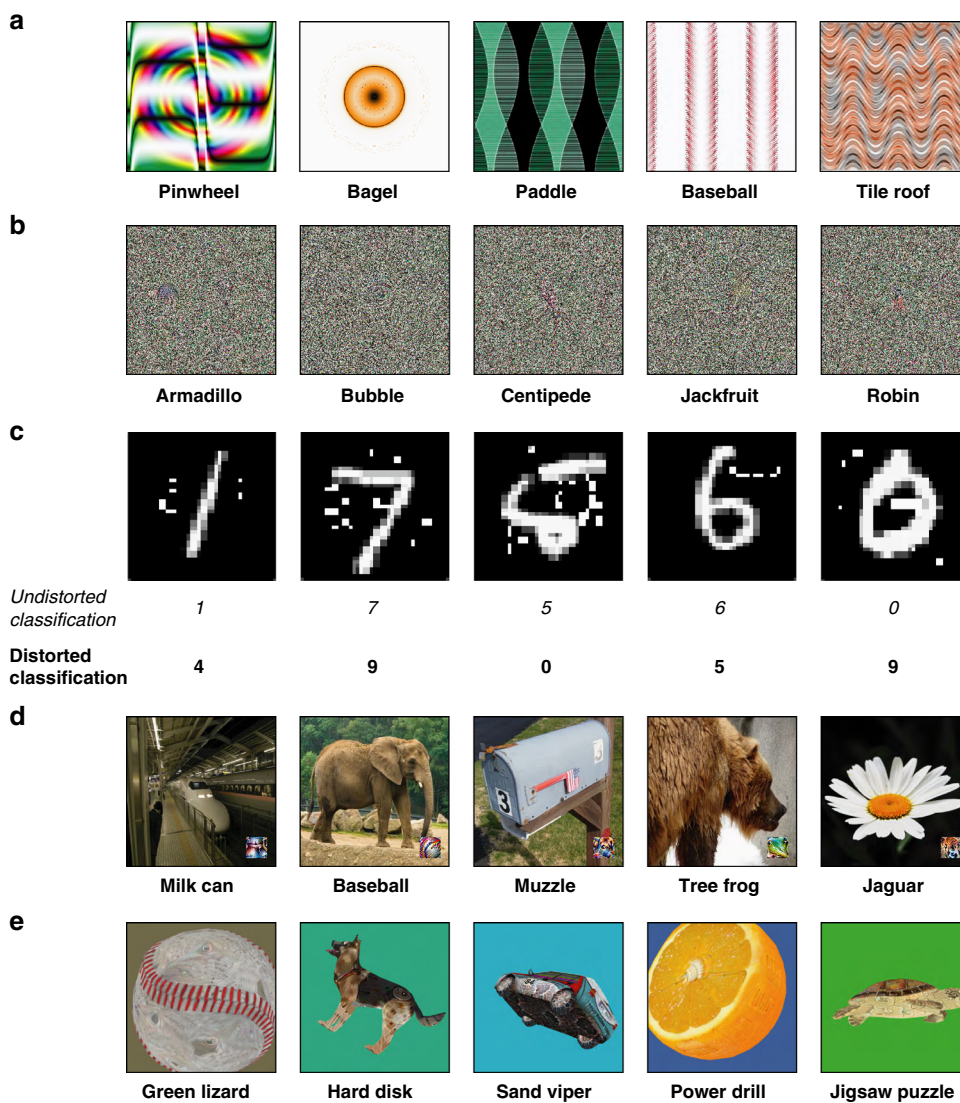


Fig. 1 Examples of adversarial images that “fool” Convolutional Neural Nets trained to classify familiar objects, with labels indicating the machine’s classifications. **a** Indirectly encoded “fooling” images (from ref. 14). **b** Directly encoded “fooling” images (from ref. 14). **c** Perturbed adversarial images that cause the machine to classify one kind of digit as another (from ref. 15). **d** The LaVAN attack¹³ can cause a machine to misclassify a natural image even when the noise is localized to a corner of the image. **e** “Robust” adversarial images¹² are renders of 3D objects that are misclassified from multiple viewpoints (and can even be physically produced in the real world)

object. So, we obtained these second-ranked choices for every fooling image, and asked observers in Experiment 2 to choose between the machine’s first choice and the machine’s second choice—i.e., between “bagel” and “pretzel” for the golden blob image, and so on for all 48 images (Fig. 2c).

Again, human observers agreed with the machine’s classifications: 91% of observers tended to choose the machine’s 1st choice over its 2nd choice, and 71% of the images showed human-machine agreement (Fig. 2d). Evidently, humans can appreciate deeper features within adversarial images that distinguish the CNN’s primary classification from closely competing alternatives. Moreover, this result also suggests that humans and machines exhibit overlap even in their rank-ordering of image labels, since Experiment 2 yielded less human-machine agreement than Experiment 1 (94% of images vs. 71% of images). This suggests that the CNN’s second-choice was also moderately intuitive to human subjects—more so than a random label, but less so than the machine’s first-choice label, just as would be expected if machine and human classification were related in this way.

Experiment 3a: Many-way classification. The above experiments show that humans can identify the machine’s preferred label from relevant alternatives. However, both of these studies involve the limited case of only two alternatives; by contrast, image-recognizing CNNs typically choose from hundreds or thousands of labels when classifying such images. Would humans exhibit reliable agreement with the machine even under more unconstrained circumstances? Although it would not be practically feasible to make humans choose from 1000 individual labels (the number of labels in ImageNet), Experiment 3 stepped closer to these conditions by displaying the labels of all 48 images at once, and asking subjects to pick the best of all the labels for each image.

Even under these demanding conditions, 88% of subjects selected the machine’s label at above-chance rates, and 79% of images showed above-chance human-machine agreement. Moreover, in an analysis inspired by the *rank-5* measure in the machine-learning literature, we found that the machine’s label was among the top five human choices for 63% of the images

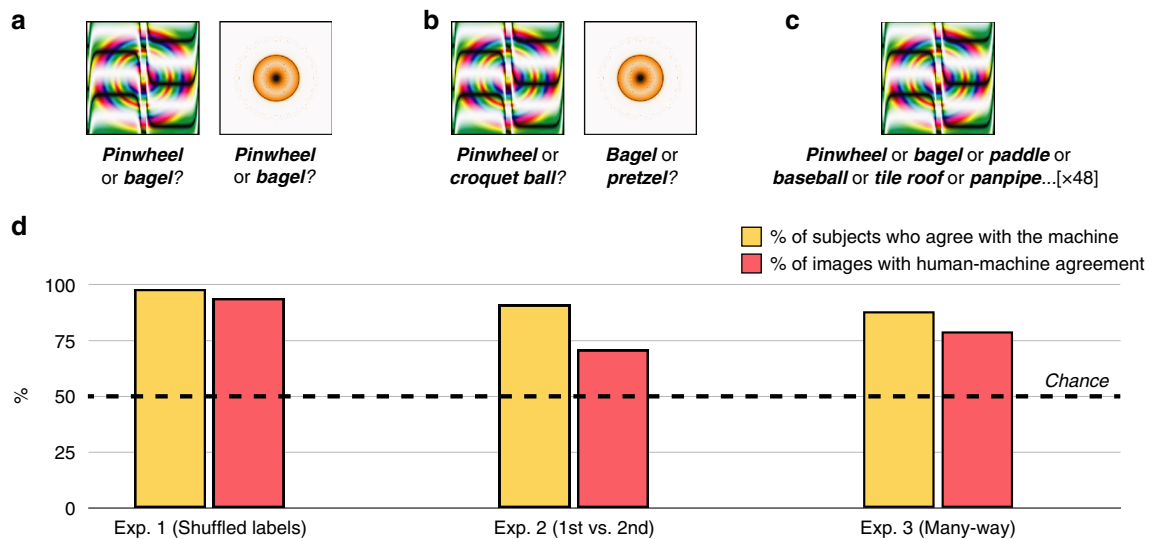


Fig. 2 Forced-choice classification with indirectly encoded fooling images. **a** In Experiment 1, 200 subjects saw one fooling image at a time (48 images total), and chose between two candidate labels: The machine’s choice for that image, and a random label drawn from the other images in the imageset. **b** In Experiment 2, 200 subjects chose between the machine’s first-choice classification and its second-choice classification. **c** In Experiment 3a, 200 subjects saw the same images as before, but with all 48 labels visible at once. **d** In all 3 experiments, most subjects agreed with the machine more often than would be predicted by chance responding (yellow bars), and most images showed human-machine agreement more often than would be predicted by chance responding (red bars). Given that raw classification accuracy in human subjects will be modulated by factors such as attention, motivation, engagement with the task, time spent studying an image, etc., we report the percentage of subjects who agreed with the machine at above-chance rates, and the percentage of images that showed above-chance human-machine agreement. For Experiment 1, the 95% confidence interval for the % of subjects with above-chance classification was [94.6% 99.4%], and it was [82.8% 98.7%] for the % of images with above-chance classification. For Experiment 2, these intervals were [87.7% 95.7%] and [58.0% 83.7%], respectively. For Experiment 3, these intervals were [83.2% 93.2%] and [67.7% 90.7%], respectively. Across all three experiments, these outcomes were reliably different from chance at $p < 0.001$ (two-sided binomial probability test)

(whereas chance responding would put this figure at 10.4%); in other words, even when the single most popular human-chosen label was not the CNN’s preferred label, the 2nd, 3rd, 4th, or 5th most popular human-chosen label (out of 48 possible choices) usually did match the CNN’s preferred label. These results suggest that humans show general agreement with the machine even in the taxing and unnatural circumstance of choosing their classification from dozens of labels displayed simultaneously.

Experiment 3b: “What is this?” The previous study more closely resembled the task faced by CNNs in classifying images, which is to classify an image by choosing one label among many. However, all of the preceding experiments differ from a CNN’s task in another way: whereas CNNs select a label that best matches an image, our human subjects were asked to anticipate the machine’s label, rather than to label the images themselves. Would humans still agree with the CNN’s classification if their task were simply to straightforwardly classify the image?

Experiment 3b investigated this question by changing the task instructions: rather than being told to “think like a machine” and guess a machine’s preferred label, subjects were simply shown images and asked “What is this?”. On each trial, an image appeared on the display, and subjects were asked “If you had to pick a label for it, what would you pick?”, from 48 possible labels. Once again, human judgments and machine classifications converged: 90% of subjects agreed with the machine at above-chance rates, and 81% of the images showed above-chance human-machine agreement. These results suggest that the humans’ ability to decipher adversarial images doesn’t depend on the peculiarities of our machine-theory-of-mind task, and that human performance reflects a more general agreement with machine (mis)classification.

Experiment 4: Television-static images. Though the images in the above experiments are peculiar, they do at least have discrete and distinguishable features; for example, the “baseball” image has a collection of cross-hatched parallel red lines that resemble the characteristic red stitching of a real baseball. (Indeed, the creators of this adversarial attack informally noted this resemblance in later work²⁶.) What about truly bizarre images that are considered “totally unrecognizable to human eyes”¹⁴?

In Experiment 4, subjects saw eight “television static” images that CNNs recognize as objects—e.g., “centipede” or “robin” (Fig. 3a). These images appear to be colorful collections of pixels with little if any underlying structure. (However, upon very close inspection, you may notice a small, often central, ‘object’ within each image.) On each trial, a given label appeared on the screen, along with five examples of that category drawn from ImageNet (e.g., the word “robin” beside five photographs of robins). Subjects were instructed to select the television-static image that best matched the label (Fig. 3a).

Even with these bizarre images, 81% of observers agreed with the machine at above-chance rates, and 100% of the images showed above-chance human-machine agreement (i.e., they were chosen as matches more than 12.5% of the time; Fig. 3c). Moreover, for 75% of the images, the label chosen most often by subjects was also the machine’s most preferred choice (analogous to *rank-1* performance). This is especially relevant for human-machine comparisons, since CNNs typically make their classification decisions after a softmax transformation has been applied to the input to the CNN’s final layer; applying a similar transformation over our human responses could thus similarly produce “high confidence” ratings for adversarial images, if we were to treat our entire human cohort’s judgments as “votes” over which a softmax decision is computed.

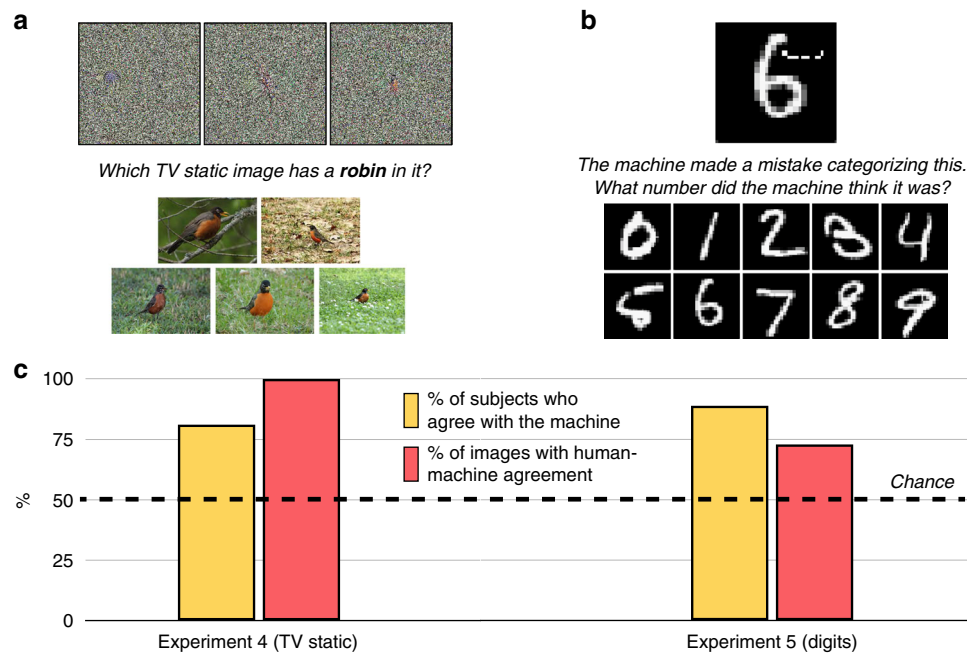


Fig. 3 Classification with directly encoded fooling images and perturbed MNIST images. **a** In Experiment 4, 200 subjects saw eight directly encoded “television static” images at once (though only three are displayed here); on each trial, a single label appeared, along with five natural photographs of the label randomly drawn from ImageNet (here, a robin [and represented in this figure using public-domain images instead of ImageNet images]). The subjects’ task was to pick whichever fooling image corresponded to the label. **b** In Experiment 5, 200 subjects saw 10 undistorted handwritten MNIST digits at once; on each trial, a single distorted MNIST digit appeared (100 images total). The subjects’ task was to pick whichever of the undistorted digits corresponded to the distorted digit (aside from its original identity). **c** Most subjects agreed with the machine more often than would be predicted by chance responding, and most images showed human-machine agreement more often than would be predicted by chance responding (including every one of the television static images). For Experiment 4, the 95% confidence interval for the % of subjects with above-chance classification was [75.5% 87.3%], and [63.1% 100%] (one-sided 97.5% confidence interval) for the % of images with above-chance classification. For Experiment 5, these intervals were [84.2% 93.8%] and [64.3% 81.7%], respectively. Across both experiments, these outcomes were reliably different from chance at $p < 0.001$ (two-sided binomial probability test)

These results suggest that human subjects are not only able to discern subtle features of adversarial images, but can also infer machine classifications of such images even when the relevant patterns are not discrete features at all but instead seemingly featureless collections of colored pixels.

Experiment 5: Perturbed digits. The foregoing experiments explored “fooling” images that human observers would not typically classify as familiar objects. However, a more insidious form of adversarial attack can occur when a few perturbed pixels fool CNNs into classifying one natural object as a different object; for example, an image that would normally be classified as a “4” might now be classified as a “7” when just a small subset of the pixels is altered¹⁵ (Fig. 1c). This sort of attack is of special practical importance: One could imagine, for example, a malicious actor altering a speed limit sign in this way, which might fool an autonomous vehicle into recognizing a *Speed Limit 45* sign as a *Speed Limit 75* sign and then dangerously accelerating as a result.

The original research that generated such images concluded that “humans cannot perceive the perturbation introduced to craft adversarial samples”¹⁵, because human observers persisted with their original classifications even after the distortion was introduced (see also ref. ²⁷). By contrast, here we asked humans which digit they would have picked if they weren’t allowed to give their initial impression. We collected 100 adversarially distorted digits that had caused a CNN (LeNet²⁸) to change its classification, and asked subjects which digit they thought the machine (mis)perceived the images as (Fig. 3b).

Even for perturbed adversarial images, human responses again aligned with the machine’s: 89% of subjects identified the machine’s classifications at above-chance rates, and 73% of images showed above-chance human-machine agreement (Fig. 3c). Thus, even when adversarial images have strong prepotent identities, humans can anticipate the machine’s misclassifications.

Experiment 6: Natural images and localized perturbations.

Whereas the previous result suggested that humans can decipher not only fooling images but also perturbed images, the particular adversarial attack explored in Experiment 5 may be limited in important ways: the proportion of perturbed pixels was often relatively high (as many as 14% of the pixels in the image); the perturbations often obstructed salient parts of the image; and the target of the adversarial attack was only handwritten digits, which differ from natural images both in their general richness and also in the breadth of possible target classes (since they involve only the digits 0–9). By contrast, more recent adversarial attacks overcome many of these limitations; could humans decipher the images produced by more advanced approaches?

Experiment 6 tested human observers on images produced by a state-of-the-art “localized” adversarial attack (“LaVAN”¹³). This adversarial distortion perturbs far fewer pixels in the attacked image; it succeeds even when the perturbation is confined to an isolated corner of the image (rather than obstructing the image’s focal object); and it can target a wider array of natural images. For example, LaVAN can cause a machine to misclassify a daisy as a jaguar or a subway train as a milk can (as in Fig. 4a), and it can do so even while perturbing only 2% of pixels near the border of the image. As the authors of that work note, however, the perturbed

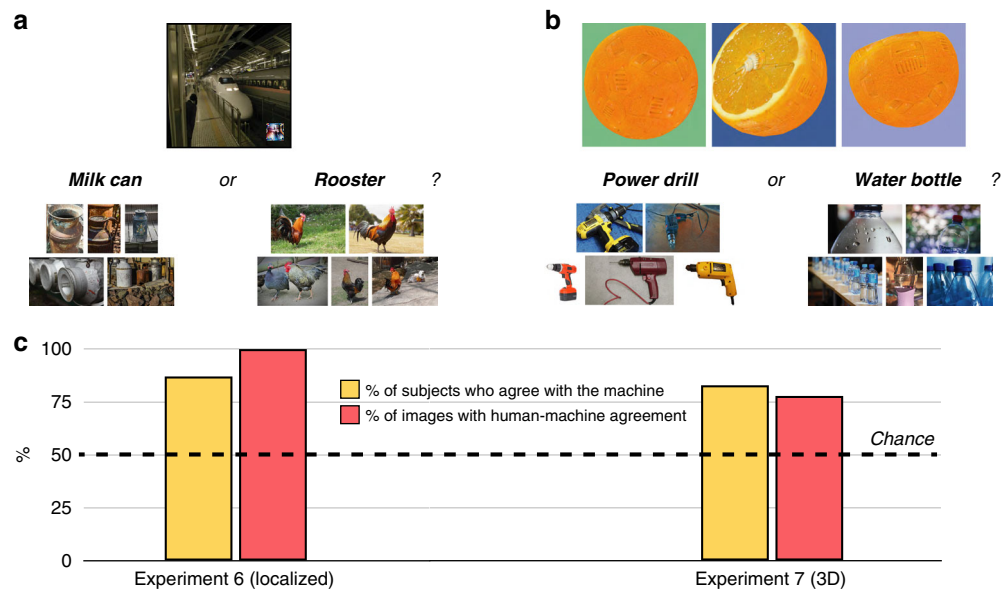


Fig. 4 Classification with perturbed natural images and 3D objects. **a** In Experiment 6, 200 subjects saw natural photographs that had a small adversarial perturbation in the bottom right corner of the image (22 images total); the subjects' task was to choose between the machine's choice for that image and a random label drawn from the other images in the imageset (which were accompanied by images of each target class for reference). **b** In Experiment 7, 400 subjects saw three viewpoints of a rendered 3D object whose adversarial perturbation caused a machine to misclassify (106 images total, of which each subject saw 53); the subjects' task was to choose between the machine's choice for that object and a random label drawn from the other objects in the imageset (with images of each target class for reference). **c** Most subjects agreed with the machine more often than would be predicted by chance responding, and most images showed human-machine agreement more often than would be predicted by chance responding (including every one of the LaVAN images). For Experiment 6, the 95% confidence interval for the % of subjects with above-chance classification was [82.5% 91.9%], and [84.6% 100%] (one-sided 97.5% confidence interval) for the % of images with above-chance classification. For Experiment 7, these intervals were [78.7% 86.5%] and [70.5% 86.1%], respectively. Across both experiments, these outcomes were reliably different from chance at $p < 0.001$ (two-sided binomial probability test)

pixels sometimes look like miniature versions of the adversarial target classes. Do naive human subjects agree?

We acquired 22 such images that caused a CNN (Inception V3²⁹) to misclassify, and placed them in the same forced-choice design as Experiment 1, but with the addition of natural images of the target and foil classes randomly taken from ImageNet (so that subjects viewing a subway-train-to-milk-can image, for example, saw the label "milk can" along with five images of milk cans drawn from ImageNet, and similarly for the foil labels; Fig. 4a). Even for this advanced adversarial attack, human responses aligned with the machine's: 87% of subjects identified the machine's classifications at above-chance rates, and 100% of the images showed above-chance human-machine agreement (Fig. 4c). Thus, even more recent and sophisticated adversarial attacks are susceptible to human deciphering.

Experiment 7: 3D objects. All of the preceding experiments explored adversarial attacks on 2D images, including both natural photographs and digitally generated textures. Such images are, certainly, the predominant targets of adversarial attacks; however, these attacks are ultimately limited in (a) their complexity, since the images are only two-dimensional; (b) their practical applications, since they typically "fool" machine classification systems only when those systems are fed such images directly in a digital format (rather than through a sensor); (c) their robustness, since most attacks on 2D images lose their fooling powers when the images are rotated, resized, blurred, or otherwise manipulated; and (d) their promise for understanding the richness of human object representation, since we typically see real-life objects in the world from multiple angles and with multiple cues, rather than a single image from a single viewpoint with only pictorial image cues.

Recently, a creative and sophisticated adversarial attack has allowed for the creation not only of adversarial 2D images but also adversarial 3D objects that machines reliably misclassify¹². When rendered using 3D graphics software, such "robust" adversarial examples continue to fool CNN-based classifiers (here, Inception V3²⁹), not only from one particular vantage point but also from multiple different angles and distances; moreover, they can even be 3D printed as physical objects in the real world. A 3D model of an orange, for example, could be produced in physical form, placed on a table, and recognized by a machine as a power drill, cucumber, or even a missile, simply because of certain vague textural elements on the orange's surface (Fig. 4b).

Experiment 7 tested human observers on such robust 3D adversarial objects. We acquired 106 such examples, encompassing 10 familiar objects (e.g., baseball, turtle, orange) that are classified as something else (e.g., lizard, puzzle, drill) when a certain texture is added to them. On each trial, human observers saw three different rendered viewpoints of each adversarial object, and were shown both the target label and a foil label drawn from another element in the imageset, with examples of the classes shown beneath (so that subjects viewing an orange-to-drill image, for example, saw the label "power drill" along with five images of power drills drawn from ImageNet, and similarly for a foil label; Fig. 4b).

Even for this state-of-the-art attack, human responses aligned with the machine's: 83% of subjects identified the machine's classifications at above-chance rates, and 78% of the images showed above-chance human-machine agreement (Fig. 4c). Once again, humans were able to decipher the machine's classifications, here for one of the most advanced and alarming adversarial attacks in the literature.

Discussion

The present results suggest that human intuition is a reliable source of information about how machines will classify images—even for adversarial images that have been specifically designed to fool the machine. This implies at least some meaningful degree of similarity in the image features that humans and machines prioritize—or can prioritize—when associating an image with a label. The very existence of adversarial images has cast into doubt whether recently developed machine-vision systems bear any real resemblance to humans in terms of how they classify images, and also whether such models can be attacked surreptitiously. The present results suggest that this particular challenge to notions of human-machine similarity may not be as simple as it appears (though there may of course be other reasons to doubt the similarity of humans and machines), and perhaps even that humans could play a role in understanding, “defending” against, or even further refining, such attacks.

We suggest here that humans can anticipate how machines will behave, even under these special and unusual circumstances. How deep does this ability run? The human subjects here showed reliable agreement with the machine across an impressively broad array of images: collages of features, television-static images, handwritten digits, natural photographs, and 3D objects. There is also reason to think that these abilities could generalize further. For example, recent work has shown that physically placing a small and colorful “sticker” next to a banana can fool CNNs into classifying images of the banana as a toaster³⁰; however, the sticker itself looks quite like a toaster, and we suspect that other attacks may be similarly decipherable.

At the same time, there is a cottage industry around the production of adversarial images, and there may well be adversarial images that humans cannot decipher in the manner explored here. For example, some kinds of adversarial images are produced by making thousands of minuscule perturbations across every pixel in the image (as in a famous panda-to-gibbon example³¹); we doubt humans could see a gibbon in this image, even under forced-choice conditions. However, there are at least two reasons why such images may actually be less powerful as challenges to human-machine comparison than the images we explored here.

First, and more practically, those examples are the very cases that exhibit the least robust transfer across systems and transformations. For example, even a small rotation or rescaling of the perturbed image is usually sufficient to return it to its prior classification, which suggests that this is not the most practical attack for real-world settings. (For example, an autonomous vehicle that photographed such images in the real world would almost certainly fail to be fooled³².) Instead, the sort of adversarial attack that is more likely to succeed against an autonomous vehicle or baggage-screening operation is exactly the sort having some sort of visible noise pattern, and so that is the sort we explored here.

Second, and more theoretically, the reason such perturbations are not visible to humans may have little to do with the high-level processes underlying human object classification, but instead with low-level physiological limitations on human visual acuity, resolution, and sensitivity to contrast, which simply cannot match the resolving power of *in silico* image processing. Moreover, whereas humans must view such images through a sensor (their eyes), machine-vision systems are typically fed such images digitally. In that case, it is plausible that humans cannot perceive or decipher panda-to-gibbon-style perturbations simply because of the limitations of their eyes and early visual systems, rather than because of the concepts or templates they deploy in classifying objects. Indeed, many instances of the panda-to-gibbon example *must* be undecipherable to humans, because the perturbation in such images is often too small to change the value of any actual pixel as

rendered on a monitor. For this reason, the mere existence of such adversarial images perhaps tells us less about similarities or differences between humans and CNNs in high-level object classification per se, but rather reflect lower-level considerations such as the resolution of human vision or even of display equipment. (Similarly, some adversarial images allow noisy pixels to take any value that the neural network can process, including those outside the dynamic range of images¹³. Humans may have difficulty deciphering those patterns as well, but perhaps not because of the principles of human object recognition.)

To be sure, our results do not suggest that adversarial images are somehow unproblematic in the applied settings for which CNNs are hoped to be useful; adversarial images remain a dangerous and alarming development. But the present results do at least suggest that human intuition about such images can be a meaningful source of information about how a machine will classify them, and even that humans could have a role to play in the “loop” that generates such images¹⁷. For example, a small minority of the images in the present experiments (e.g., 3/48 in Experiment 1) had CNN-generated labels that were actively rejected by human subjects, who failed to pick the CNN’s chosen label even compared to a random label drawn from the imageset. Such images better meet the ideal of an adversarial example, since the human subject actively rejects the CNN’s label. However, we note that it was not clear in advance of collecting the human data exactly which images the humans would be able to decipher and which they would not. An important question for future work will be whether adversarial attacks can ever be refined to produce only those images that humans cannot decipher, or whether such attacks will always output a mix of human-classifiable and human-unclassifiable images; it may well be that human validation will always be required to produce such truly adversarial images (and that human testing on candidate adversarial images should be incorporated into the pipeline of testing and validating new CNN-based models of object classification).

Indeed, one could state this possibility as a conjecture (call it the “knowable noise” conjecture): As long as (a) an adversarial attack produces noise that is visible to a physical sensor (like our eyes), (b) the adversarial image is robust and transferable, and (c) the to-be-attacked system demonstrates human-level recognition accuracy on a wide array of images, that attack will tend to produce images that are judged by humans to resemble their target class.

A related question is whether human subjects could, with training, improve their ability to decipher adversarial images. For example, Experiment 4 (with “television static” images) involved perhaps the most challenging and unfamiliar sorts of adversarial images, and exhibited a powerful practice effect, with a strong positive correlation between trial number and classification accuracy, evident even as a simple linear correlation, $r(6) = 0.79$, $p < 0.02$. (This can also be shown by comparing classification accuracy on the final trial vs. the first trial, $t(166) = 3.19$, $p < 0.002$.) This suggests that greater familiarity with the space of adversarial images might allow humans to better anticipate the machine’s classifications, and perhaps that future work could determine how best to prepare and train humans to detect and decipher such images.

What do these results say about the relationship between humans and machines? An important property of the adversarial examples studied here is that they were originally created without the human visual system in mind. Other work has produced images that cause humans to misclassify under choice- and time-limited circumstances (e.g., classifying an image of a distorted “dog” as a “cat” when the image is presented for 63 ms²⁴). The conclusions of this work are consonant with our own, in that they show how humans and CNNs can be made to give similar

classifications for adversarial images. However, one important difference is that the success of this earlier work required explicitly incorporating aspects of the human image-processing stream into the procedure for generating adversarial images. For example, the adversarial images produced by that procedure not only had to fool a CNN into misclassifying, but also had to first pass through models of the human retina and sophisticated forms of spatial blurring that incorporate real measurements from the primate visual system. By contrast, the images explored in the present studies were simply generated to fool a machine. In at least this sense, the present studies should have “stacked the deck” against human-machine convergence, since the adversarial images we study here were generated without any consideration of human vision at all; yet, we still find evidence for human deciphering of adversarial stimuli.

How, then, did our human subjects do this? As alluded to earlier, the answer may be in part that adversarial examples truly do share core visual features with the images they are mistaken for, especially considering the available labels. (Indeed, this simple fact may help explain why adversarial images generated for one CNN often transfer to other CNNs²⁵.) Why, then, does it seem so strange that such images should be classified as familiar objects? To be sure, it is unlikely that subjects in our experiments truly recognized most of the adversarial images shown to them, in the sense of rapidly and spontaneously matching the image to a stored object representation; for example, it seems unlikely that humans could easily identify the adversarial images’ target classes without at least some idea of the relevant label options. However, this possibility does not undermine the interest of the present results, for at least three reasons.

First, even in the absence of spontaneous recognition, humans can engage in surprisingly sophisticated processing of even very sparse textures³³, and object identification in humans benefits in important and well-established ways from image labels and other contextual factors³⁴—especially for ambiguous or degraded images³⁵. For this reason, it is only natural that explicit labels and examples assist our human subjects, and that a given object identity isn’t immediately forthcoming upon looking at the adversarial images. (Indeed, machine-vision systems themselves also don’t engage in “free classification” when they process adversarial images; they simply pick the best label in their provided vocabulary, just as our human subjects did.) Second, the real-world situations in which humans might one day encounter adversarial images may themselves involve known constraints on the relevant target classes: for example, if a human sees some odd patterns on the number in a speed limit sign and suspects it may be an adversarial image, the space of possible target classes may be fairly limited (as in Experiment 5). Third, the distinctions made in cognitive science between rapid, effortless recognition and slower, more deliberate reasoning simply do not exist for CNNs, whose architectures cannot easily be parsed in these ways. Though this very fact suggests an even deeper difference between humans and CNNs (for classic critiques of similar approaches, see^{36,37}), it also means that we cannot be sure that today’s CNNs are doing genuine recognition either. In other words, even though our experiments may tap into human “cognition” more than human “perception”³⁸, these distinctions may not even exist for CNNs—and so both the CNNs’ behavior and the humans’ behavior might be readily interpreted as simply playing along with picking whichever label is most appropriate for an image.

Indeed, although adversarial images are often analogized to optical illusions that flummox human vision^{7,11,39}, we suggest another analogy: Whereas humans have separate concepts for appearing *like* something vs. appearing *to be* that thing—as when a cloud looks like a dog without looking like it *is* a dog, or a snakeskin shoe resembles a snake’s features without appearing to

be a snake, or even a rubber duck shares appearances with the real thing without being confusable for a duck—CNNs are not permitted to make this distinction, instead being forced to play the game of picking whichever label in their repertoire best matches an image (as were the humans in our experiments). After all, the images in Fig. 2a do look like pinwheels and bagels (at least, more than they look like baseballs or roofs); they just don’t look like they *are* pinwheels and bagels. Perhaps CNNs would agree, if they could.

Methods

General Methods for Experiments 1-7. Participants: In all of the experiments reported here, separate groups of 200 subjects participated online through Amazon Mechanical Turk (for validation of this subject pool’s reliability, see ref. 40). (In Experiment 7, 400 subjects participated, being randomly assigned to see one or another half of the images). All groups of subjects (1800 total) provided informed consent and were compensated financially for their participation.

Procedure: For the machine-theory-of-mind task, subjects were told about “a machine that can look at a picture and tell us what it is”, and also that the machine sometimes “gives surprising answers”. Subjects were told that the images that cause the surprising answers were collected here, and that their job was to guess what answer the machine gave (except in Experiment 3b, in which subjects were simply asked to classify the images). In Experiments 1, 2, 3a, 3b, 5, 6, and 7, adversarial images were displayed in different random orders for each subject (with each subject seeing each image exactly one time, except in Experiment 7, where each subject saw half of the images exactly one time each), and subjects clicked a button to indicate which label they thought the machine gave. In Experiment 4, all 8 adversarial images were visible on every trial, and instead various candidate labels appeared in different random orders for each subject, with subjects picking the adversarial image that best matched the label shown on that trial. The response options also appeared in random locations, with the machine’s “true” answer being equally likely to appear in any button location. After giving a response, the images and buttons disappeared for 500 ms, after which the next trial appeared. In all experiments, subjects who quit early or otherwise failed to submit a complete dataset were excluded from further analysis, as were subjects whose median response time across all trials was less than 1000 ms, which suggested that they simply clicked through the experiment without actively participating. Post-exclusion sample sizes for Experiments 1–7 were 185 (E1), 181 (E2), 161 (E3a), 174 (E3b), 167 (E4), 164 (E5), 195 (E6), and 368 (E7). However, no result reported here depended in any way on these exclusions; i.e., every pattern remained statistically reliable even without excluding any subjects.

Adversarial stimuli: Experiments 1–3 used a set of 48 indirectly encoded “fooling” images obtained from ref. 14; Experiment 4 used 8 additional directly encoded images. Experiment 5 used 100 distorted images appearing in ref. 15. Experiment 6 used 22 distorted images generated by ref. 13 that resulted in a >75%-confidence classification of the adversarial target class. Experiment 7 used 106 images generated by¹² that resulted in a >95%-confidence classification of the adversarial target class from at least 3 views (with the top 3 such views displayed to subjects on a given trial). All images appeared to subjects at their native resolution, without any additional compression or distortion (unless subjects actively zoomed their browser in or out).

Data availability

All data, code, and materials that support the findings of this study are available at <https://osf.io/uknbn>.

Received: 27 October 2018 Accepted: 21 January 2019

Published online: 22 March 2019

References

1. Krizhevsky, A., Sutskever, I. & Hinton, G. E. Imagenet classification with deep convolutional neural networks. *Adv. Neural Inf. Process. Syst.* **25**, 1097–1105 (2012).
2. Russakovsky, O. et al. ImageNet large scale visual recognition challenge. *Int. J. Comput. Vis.* **115**, 211–252 (2015).
3. LeCun, Y., Bengio, Y. & Hinton, G. Deep learning. *Nature* **521**, 436–444 (2015).
4. Cichy, R. M., Khosla, A., Pantazis, D., Torralba, A. & Oliva, A. Comparison of deep neural networks to spatio-temporal cortical dynamics of human visual object recognition reveals hierarchical correspondence. *Sci. Rep.* **6**, 27755 (2016).

5. Greene, M. R. & Hansen, B. C. Shared spatiotemporal category representations in biological and artificial deep neural networks. *PLoS. Comput. Biol.* **14**, e1006327 (2018).
6. Jozwik, K. M., Kriegeskorte, N., Storrs, K. R. & Mur, M. Deep convolutional neural networks outperform feature-based but not categorical models in explaining object similarity judgments. *Front. Psychol.* **8**, 1726 (2017).
7. Kriegeskorte, N. Deep neural networks: A new framework for modeling biological vision and brain information processing. *Annu. Rev. Vis. Sci.* **1**, 417–446 (2015).
8. Kubilius, J., Bracci, S. & de Beeck, H. P. O. Deep neural networks as a computational model for human shape sensitivity. *PLoS. Comput. Biol.* **12**, e1004896 (2016).
9. O'Connell, T. P. & Chun, M. M. Predicting eye movement patterns from fMRI responses to natural scenes. *Nat. Commun.* **9**, 5159 (2018).
10. Peterson, J. C., Abbott, J. T., & Griffiths, T. L. Adapting deep network features to capture psychological representations. In *Proc. of the 48th Annual Conference of the Cognitive Science Society*, 2363–2368 (2016).
11. Yamins, D. L. & DiCarlo, J. J. Using goal-driven deep learning models to understand sensory cortex. *Nat. Neurosci.* **19**, 356–365 (2016).
12. Athalye, A., Engstrom, L., Ilyas, A., & Kwok, K. Synthesizing robust adversarial examples. In *Proc. of the 35th International Conference on Machine Learning*, 284–293 (2018).
13. Karmon, D., Zoran, D., & Goldberg, Y. LaVAN: Localized and visible adversarial noise. In *Proc. of the 35th International Conference on Machine Learning*, 2507–2515 (2018).
14. Nguyen, A., Yosinski, J., & Clune, J. Deep neural networks are easily fooled: High confidence predictions for unrecognizable images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 427–436 (2015).
15. Papernot, N., McDaniel, P., Jha, S., Fredrikson, M., Celik, Z. B., & Swami, A. The limitations of deep learning in adversarial settings. In *IEEE European Symposium on Security and Privacy*, 372–387 (2016).
16. Szegedy, C., et al. Intriguing properties of neural networks. In *Proc. ICLR* (2014).
17. Biggio, B. & Roli, F. Wild patterns: ten years after the rise of adversarial machine learning. *Pattern Recognit.* **84**, 317–331 (2018).
18. Brendel, W., et al. Adversarial vision challenge. *Preprint at <https://arxiv.org/abs/1808.01976>* (2018).
19. Eykholt, K., et al. Robust physical-world attacks on deep learning visual classification. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition*, 1625–1634 (IEEE, New Jersey, 2018).
20. Griffiths, T. L., Abbott, J. T. & Hsu, A. S. Exploring human cognition using large image databases. *Top. Cogn. Sci.* **8**, 569–588 (2016).
21. Guo, Y. et al. Deep learning for visual understanding: A review. *Neurocomputing* **187**, 27–48 (2016).
22. Rajalingham, R. et al. Large-scale, high-resolution comparison of the core visual object recognition behavior of humans, monkeys, and state-of-the-art deep artificial neural networks. *J. Neurosci.* **38**, 7255–7269 (2018).
23. Chandrasekaran, A., Yadav, D., Chattopadhyay, P., Prabhu, V., & Parikh, D. It takes two to tango: Towards theory of AI's mind. *Preprint at <https://arxiv.org/abs/1704.00717>* (2017).
24. Elsayed, G. F. et al. Adversarial examples that fool both human and computer vision. *Adv. Neural Inf. Process. Syst.* **31**, 3914–3924 (2018).
25. Tramèr, F., Papernot, N., Goodfellow, I., Boneh, D., & McDaniel, P. The space of transferable adversarial examples. (2017) <https://arxiv.org/abs/1704.03453>.
26. Nguyen, A., Yosinski, J. & Clune, J. Understanding innovation engines: Automated creativity and improved stochastic optimization via deep learning. *Evol. Comput.* **24**, 545–572 (2016).
27. Harding, S. M., Rajivan, P., Bertenthal, B. I., & Gonzalez, C. Human decisions on targeted and non-targeted adversarial samples. In *Proc. of the 40th Annual Conference of the Cognitive Science Society* (eds. Rogers, T. T., Rau, M., Zhu, X. & Kalish, C. W.) 451–456 (2018).
28. LeCun, Y., Bottou, L., Bengio, Y., & Haffner, P. Gradient-based learning applied to document recognition. In *Proc. IEEE* 2278–2324 (1998).
29. Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., & Wojna, Z., Rethinking the Inception architecture for computer vision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* 2818–2826 (2016).
30. Brown, T. B., Man, D., Roy, A., Abadi, M., & Gilmer, J. Adversarial patch. In *Proc. 31st Conference on Neural Information Processing Systems (NIPS)*, (2017).
31. Goodfellow, I. J., Shlens, J., & Szegedy, C. Explaining and harnessing adversarial examples. In *Proc. ICLR*. (2015).
32. Lu, J., Sibai, H., Fabry, E., & Forsyth, D. No need to worry about adversarial examples in object detection in autonomous vehicles. *Preprint at <https://arxiv.org/abs/1707.03501>* (2017).
33. Long, B., Störmer, V. S. & Alvarez, G. A. Mid-level perceptual features contain early cues to animacy. *J. Vis.* **17**, 20–20 (2017).
34. Lupyan, G. & Thompson-Schill, S. L. The evocative power of words: activation of concepts by verbal and nonverbal means. *J. Exp. Psychol. Gen.* **141**, 170–186 (2012).
35. Bar, M. Visual objects in context. *Nat. Rev. Neurosci.* **5**, 617–629 (2004).
36. Fodor, J. A. & Pylyshyn, Z. W. Connectionism and cognitive architecture: A critical analysis. *Cognition* **28**, 3–71 (1988).
37. Marcus, G. F. Rethinking eliminative connectionism. *Cogn. Psychol.* **37**, 243–282 (1998).
38. Firestone, C. & Scholl, B. J. Cognition does not affect perception: evaluating the evidence for 'top-down' effects. *Behav. Brain Sci.* **e229**, 1–77 (2016).
39. Majaj, N. J. & Pelli, D. G. Deep learning Using machine learning to study biological vision. *J. Vis.* **18**, 1–13 (2018).
40. Crump, M. J. C., McDonnell, J. V. & Gureckis, T. M. Evaluating Amazon's Mechanical Turk as a tool for experimental behavioral research. *PLoS ONE* **8**, e57410 (2013).

Acknowledgements

For helpful correspondence and for sharing images and details from their work, we thank Anish Athalye, Danny Karmon, Anh Nguyen, and Nicolas Papernot. For helpful discussion and/or comments on earlier drafts, we thank Wieland Brendel, Christopher Honey, Tom Wallis, Brad Wyble, and Daniel Yamins. Z.Z. was supported by the JHU Office of Undergraduate Research. C.F. was supported by the JHU Science of Learning Institute.

Author contributions

C.F. and Z.Z. contributed to the design and execution of the experiments and data analyses, and wrote the paper in collaboration.

Additional information

Competing interests: The authors declare no competing interests.

Reprints and permission information is available online at <http://npg.nature.com/reprintsandpermissions/>

Journal peer review information: *Nature Communications* thanks the anonymous reviewers for their contributions to the peer review of this work.

Publisher's note: Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2019