arXiv:1903.01069v2 [cs.LG] 5 Mar 2019

Do Neural Networks Show Gestalt Phenomena? An Exploration of the Law of Closure

Been Kim [1] **Emily Reif** [1] **Martin Wattenberg** [1] **Samy Bengio** [1]
beenkim@google.com, Google

Abstract

One characteristic of human visual perception is the presence of "Gestalt phenomena," that is, that the whole is something other than the sum of its parts. A natural question is whether image-recognition networks show similar effects. Our paper investigates one particular type of Gestalt phenomenon, the law of closure, in the context of a feedforward image classification neural network (NN). This is a robust effect in human perception, but experiments typically rely on measurements (e.g., reaction time) that are not available for artificial neural nets. We describe a protocol for identifying closure effect in NNs, and report on the results of experiments with simple visual stimuli. Our findings suggest that NNs trained with natural images do exhibit closure, in contrast to networks with randomized weights or networks that have been trained on visually random data. Furthermore, the closure effect reflects something beyond good feature extraction; it is correlated with the network's higher layer features and ability to generalize.

1. Introduction

The field of cognitive psychology is devoted to understanding a black box system of interest – humans. Its progress includes the discovery of fundamental rules of how humans operate (Schultz & Schultz, 2015). For example, one overarching set of rules, proposed in 1923 to explain many different perceptual phenomena, is known as 'Gestalt' laws (Wertheimer, 1923). These principles have had a huge impact on modern psychology (Kimchi, 1992; Wagemans et al., 2012a;b; Schultz & Schultz, 2015) despite some criticism (Schultz & Schultz, 2015; Wagemans et al., 2012a;b).

We are interested in studying another black box of interest – neural networks. Given the successes of cognitive psychology, and the potential parallels between biological and artificial neural networks (Cadieu et al., 2014; Riesenhuber & Poggio, 1999; Ritter et al., 2017), it is natural to look to psychology to gain insight. At the same time, knowledge of similarities and differences between human and NN's perception is helpful for interpretability, especially how humans interpret explanations. When presented with an 'explanation', humans' interpretation is fundamentally constrained by the way we perceive the world, because we have never seen it any other way. For instance, knowing that some things that are obvious to us do not hold in NNs extends our capacity to better create and understand explanations.

In this paper, we focus on one particular Gestalt principle, the *Law of Closure* (Wertheimer, 1923). This law asserts that the human visual perception system has a tendency to "close the gap" in order to perceive whole objects[1] when only fragments are visible – for instance, to see a complete triangle even when the sides have missing pieces (e.g., Fig. 2). Some have argued this tendency stems from the mind's natural tendency to recognize familiar patterns, filling in any information that may be missing – or in more statistical language, from learning the underlying distribution of image data (Kimchi, 1992).

We aim to determine whether NNs exhibit similar closure effects. That is, are the shapes in set Illusory in Fig. 2 perceived as complete shapes, or simply a collection of lines? The experimental framework typically used for humans requires measurements (e.g., reaction time) that are not usually available for artificial NNs. Here, we propose a simple protocol for measuring closure effects in NNs.

We train simple and complex NNs with natural images and test them with simple visual illusory stimuli. The results suggest that trained image classification systems do show closure effects, in contrast to networks with randomized weights or that have been trained on visually random data. More interestingly, the closure goes beyond the NN's feature extraction ability, and is correlated to the network's ability to generalize. Analogous to human studies (von der Heydt et al., 1984), higher layers that relate to higher level features typically show closure effects more than earlier layers.

Contributions:

[1] What an object is in visual perception is a difficult question to answer (Feldman, 2003). We follow to Kimchi et al. (2016)

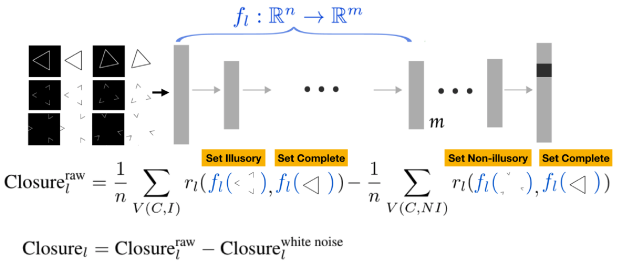

Figure 1. The experiment framework for measuring closure effect in NN

• We propose an experimental framework to test the closure effect in NNs.
• We present a set of experiments which suggest that the closure effect goes beyond the NN's feature extraction ability, and is correlated to the network's ability to generalize.
• These experiments highlight potential similarities and differences between closure effects in NN and humans.

2. Related work

The key idea of the Gestalt psychology school is that we perceive individual sensory stimuli as meaningful wholes (Wertheimer, 1923). Further, the Gestalt psychologists maintained that when sensory elements are combined, the elements form a new pattern or configuration. For example, when you hear musical notes, a melody emerges from their combinations, something that did not exist in individual elements. In other words, the whole is different from the sum of its parts (Koffka, 2013). This overarching idea explains many phenomena of human perception (Wertheimer, 1923), one of which is illusory contours (set Illusory in Fig. 2); the brain has a need to see familiar simple objects and has a tendency to create a "whole" image from individual elements.

Although Gestalt psychology has faced some criticism over a lack of rigor (Wagemans et al., 2012a; Westheimer, 1999; Schultz & Schultz, 2015), investigators have successfully operationalized several of its key concepts in psychology (Ren & Malik, 2003), and it has influenced work in medicine (Bender, 1938), computer vision (Desolneux et al., 2007), therapy (Zinker, 1977), and design (Behrens, 1998).

This work is also closely related to investigations of the representations learned by NNs, an area of active study. Particularly relevant are studies comparing the representations of NN to humans' neural responses. Some have found that the representation and behavior of deep convolutional neural networks are similar to primate visual systems (Cadieu et al., 2014; Cichy et al., 2016; Yamins et al., 2014), while others found a model of object recognition in the human cortex that is similar to some aspects of modern NNs (Riesenhuber & Poggio, 1999). The different types of behavior and mistakes of humans and NN are also studied, as the conditions of classification task change (Geirhos et al., 2017; Rajalingham et al., 2018; Eckstein et al., 2017), and when confronted with adversarial examples (Elsayed et al., 2018).

Leveraging the experimental framework of classical psychology to study NNs is an important yet under-explored area of research. Recent work in this direction includes investigating shape biases in NNs (Ritter et al., 2017) or measuring abstract reasoning ability using tests designed for humans (Barrett et al., 2018). Our work aims to continue this effort.

3. Measuring closure effects and sanity check

We present our method of measuring closure effects, which borrows ideas from standard human experimental paradigms. We also describe the results of a simple sanity check on this technique.

3.1. A review of human experimental paradigms

Gestalt phenomena have been tested in humans in many ways: by measuring reaction time (Kimchi et al., 2016), discrimination ability (Ringach & Shapley, 1996), analyzing EEG data (Sanguinetti et al., 2016), building ideal observer models (Machilsen et al., 2016; Erlikhman & Kellman, 2016) and many others (for benefits and criticisms of each approach, see (Jkel et al., 2016)). For example, Kimchi et al. (2016) used subjects' response time to study relationship between perceptual organization and attention, while Ringach & Shapley (1996) measured performance on discrimination (i.e., classification) tasks to study factors that matter for illusory response.

Measuring closure effects in feedforward classification networks is not straightforward, since none of the human experimental paradigms can be directly used. While Elsayed et al. (2018) proposed a new experimental setup to accounting for the fact that time cannot be used to study neural networks, we consider a simpler framework in this paper. Ideal Observer model analysis (Geisler, 1989) is promising, but its assumptions, as well as the right way to build this model, are much-debated (Ma, 2012).

3.2. Neural network experimental paradigm

In this work, we propose a simple technique to measure closure effects in NNs (Fig. 1). As with analyses of EEG data, our method attempts to find the difference between the network's response to a "whole" figure stimulus versus a fragmentary stimulus using activations in intermediate layers of the network.

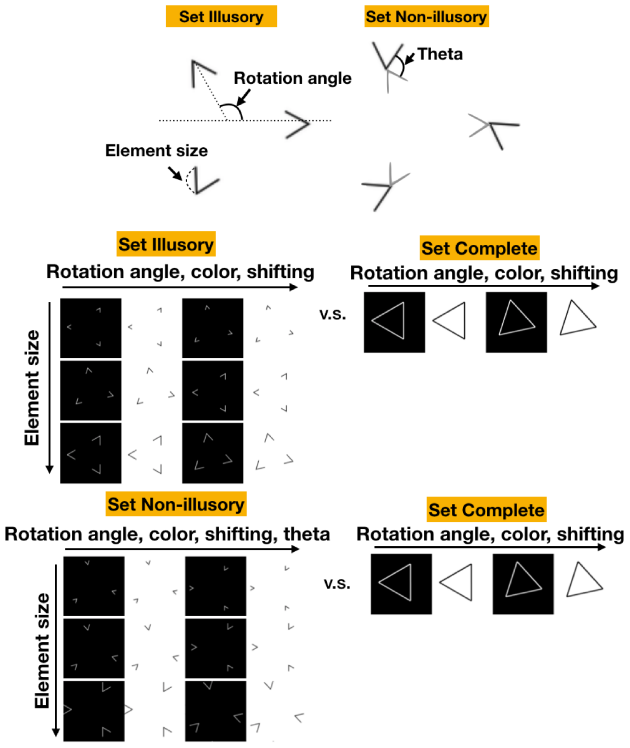

Figure 2. Examples of visual stimuli used. Element size represents the amount of space removed from the complete shape. As element size increases, it becomes easier to see a full triangle. A stimulus has black or white background color.

3.2.1. Sets of Illusory and Non-Illusory Stimuli

We use three sets of artificial stimuli, similar to those used in classic experiments with humans (e.g., Elder & Zucker (1993); Kimchi et al. (2016); Kimchi (1992)) as shown in Fig. 2. These allow us to model our experiments after existing techniques in cognitive psychology. We also found, in initial explorations with more naturalistic images, that it was difficult to control for various potential confounds (e.g., other Gestalt effect than closure). The three sets include:

- **Complete** (C): Complete triangles.
- **Illusory** (I): Triangles where sections of edges have been removed; these are still perceived as "wholes" by humans due to the closure effect.
- **Non-illusory** (NI): Images similar to those in I, except that the remaining triangle corners have been rotated (by angle θ) in order so that the resulting image is no longer seen as a whole (following Kimchi et al. (2016)).

A key independent variable in similar experiments on humans is how much of the triangles in set I remain visible (Elder & Zucker, 1993; Jkel et al., 2016). The larger the remaining "corners" in the image, the stronger closure effects are in humans. We refer to this variable as **element size**, and varied it systematically in both the I and NI sets.

To avoid confounds based on superficial image similarities (e.g., L^2 distance), and to create a data set large enough to have statistical power to measure small effects, we varied four other variables: background color (either black or white), position of images (ranging over small pixel translations), rotation of whole image, and for NI, values of θ (always at least fifteen degrees). We did not vary the scale of the images, following Ringach & Shapley (1996)'s findings that the illusion is scale invariant to humans. These variations led to 992 distinct stimuli images.

3.2.2. Measuring Closure Effects

At a high level, we compare a NN's activation similarities between the three sets of stimuli in Fig. 2. Our discrimination task (similar to Elder & Zucker (1993)) is comparing similarities between the network's responses to set I and set C with similarities between set NI and set C. If responses to set C are more similar to set I than the set NI, it suggests the presence of a closure effect.

To make the notion of "similarity" precise, consider a neural network model with inputs $\boldsymbol{x} \in \mathbb{R}^n$ and a feedforward layer l with m neurons, such that input inference and its layer l activations can be seen as a function $f_l : \mathbb{R}^n \to \mathbb{R}^m$.

Raw closure measurements For two inputs x and y, we define the *response similarity* in layer l as the cosine similarity between the activations during inference:

$$r_l(x, y) = \frac{f_l(x) \cdot f_l(y)}{|f_l(x)||f_l(y)|} .$$

To measure the aggregate similarity between responses to images in C, I, and NI, we average the responses between pairs, but with one twist: we discard pairs of images if both images in the pair have the same rotation angle, and therefore may have a spuriously small L^2 distance. We do this to minimize potential confounding reflecting pixel-level similarity. We call the remaining pairs of images *valid* pairs, and denote by $V(C, I)$ and $V(C, NI)$ the sets of valid pairs for comparing between the respective sets. Both sets have the same size, which we call n.

Comparing average differences between valid pairs leads to what we call a *raw* closure measure:

$$\text{Closure}_l^{\text{raw}} = \frac{1}{n} \sum_{V(C,I)} r_l(x, y) - \frac{1}{n} \sum_{V(C,NI)} r_l(x, y)$$

Note that one can use other similarity measures, such as canonical correlation analysis (Härdle & Simar, 2007) or representational similarity analysis (Kriegeskorte et al., 2008) (mostly designed to compare two signals with different dimensions, e.g., brain-activity measurement and behavioral measurement).

Baseline correction As our framework resembles similar human experiments using EEG analysis, we perform baseline correction, as commonly done in analyzing such high dimensional signals (Keil et al., 1999; Luck, 2014). For example, in human experiments, neural responses in between presentations of stimuli are commonly used as baselines (Grandchamp & Delorme, 2011).

To set a baseline for a given image, we use the activations that result from running it through a network trained with white noise images (WHITE NOISE condition). A plausible alternative would have been to use the UNTRAINED network's response, and indeed this led to similar results in early tests. However, we chose WHITE NOISE, since results are less dependent on the precise choice of initialization method. (And speculatively, the WHITE NOISE network may also be more analogous to EEG data when humans subjects are not engaged in any activities.)

We then use this baseline to adjust the raw closure values in aggregate, averaging over all stimuli for WHITE NOISE, rather than per-stimulus-pair. [2] That is, we adjust the closure effect measurement for a given network is by subtracting off the raw closure effect from WHITE NOISE.

Our final measurement is:

$$\text{Closure}_l = \text{Closure}_l^{\text{raw}} - \text{Closure}_l^{\text{white noise}}$$

We also report Welch's t-test (Welch, 1947) with individual samples of $r_l(x, y)$ for $V(C, I)$ and $r_l(x, y)$ for $V(C, NI)$. The null hypothesis is whether the mean of the two distributions, $r_l(x, y)$ for $V(C, I)$ and $r_l(x, y)$ for $V(C, NI)$ are the same, assuming unequal variance.

3.3. Closure measurement method sanity check

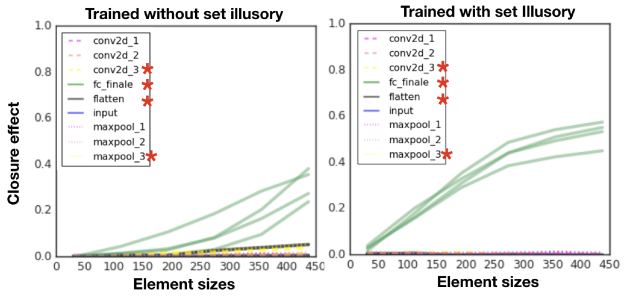

Figure 3. Sanity check experiment: Networks trained with set I have more closure effect. Multiple lines represent multiple runs of the same condition.

[2] Whether to adjust at an aggregate level or per stimulus is a much-debated question in working with EEG data (Keren et al., 2010; Plöchl et al., 2012; Maess et al., 2016; Tanner et al., 2016). Alternative baseline methods for NNs would be a good area for further studies.

We performed a simple "sanity check" on our closure effect measurement technique, comparing a network which has been trained to see sets C and I as similar, with a network which has not. Our aim in this experiment is not to explore whether closure effects exist in naturally occurring systems, but rather to test our measuring framework: if we can detect closure effects anywhere, it should find them here.

The training is as follows: both networks are trained to classify triangles from non-triangles (set Non-illusory). The first network is trained to distinguish set Complete, C from set Non-illusory, NI. The second network is trained in the same way, except that the positive examples include the set Illusory, I, as well as the set Complete. Both trained networks achieved very small validation error.

The results of this experiment show that different layers see different levels of closure effects, which we explore this in detail below. However, the main effect is that–as expected–the technique detects a large difference between the two networks, with the closure effect seen more strongly in the second network (Fig.3), reflecting the "ground truth" of the set-up. Note that with the largest element sizes, however, the first network does show what appears to be a closure effect as well; we speculate this is because for large element sizes the stimuli in C and I are far more visually similar than C and NI.

4. Hypotheses and experimental setup

The high-level goal of our experiments is to tease out important factors for closure effects in NNs. We first lay out a set of tested hypotheses and provide details of our experiments.

4.1. Hypotheses

We formulated several hypotheses on what factors might be associated with stronger closure effects:

H1. The closure effect is associated with generalization.
H2. The closure effect is stronger in higher layers than lower layers.
H3. The closure effect will generally increase during training before convergence.
H4. The closure effect is NOT arbitrarily influenced by simple input manipulations (e.g., brightness).
H5. The closure effect is stronger in deeper networks.
H6. The closure effect is stronger when trained with intentionally occluded images.
H7. The closure effect is stronger with convolutional operators than without.

We train many NNs with different running conditions described in Fig. 4 to prove or disprove the above hypotheses. For example, by comparing BREGMAN OCCLUSION and BAR OCCLUSION to NORMAL networks, we can see how

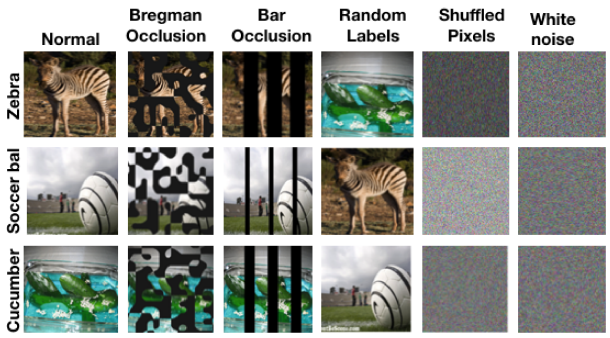

Type	Trained
NORMAL	with 600 images for each nc classes
BREGMAN OCCLUSION	with images occluded by structured noise patterns (Bregman, 2017)
BAR OCCLUSION	with images occluded by vertical black bars
RANDOM LABELS nc	with randomly labeled images of nc classes
RANDOM LABELS 1000	with randomly labeled images of 1000 classes
SHUFFLED PIXELS	with images of nc classes. Pixels are shuffled across channels.
WHITE NOISE	with random white noise images.
UNTRAINED	is an untrained network.
SMALL DATA	with one image for each nc classes.

Figure 4. Descriptions of running conditions and a subset of training data used to test hypotheses

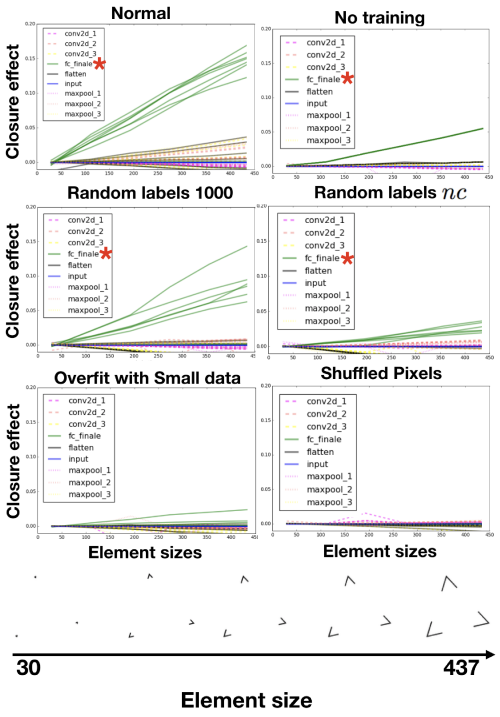

Figure 5. **H1.** Each graphs shows the closure effect as element size increases for each layer for each running condition. Multiple lines represent multiple runs for the same condition. Note that input (blue line) stays very small as we control the pixel overlap. Layers that resulted in statistical significance are marked with red $*$.

much occluded images matter to the effect. The RANDOM LABELS nc and RANDOM LABELS 1000 networks may learn representations (i.e., training error is close to zero), but they are not able to generalize (i.e., test error is big). RANDOM LABELS 1000 has access to more variety of images to extract features from than RANDOM LABELS nc. SHUFFLED PIXELS networks cannot learn good representations, nor can they generalize. However, they are different from WHITE NOISE networks, as they have access to some statistics of natural images (e.g., averaged pixel values). We train each condition with convolutional layers and fully-connected layer only (FC-only) networks.

4.2. Experimental setup

We first train a number of simple networks with which we can iterate quickly. Then, we extend our findings to a more complex, widely used network (Szegedy et al., 2016). We test the closure effect at a wide range of element sizes, and at each layer of the network.

4.2.1. SIMPLE NETWORK

Each network has $nc \in \{3, 6, 9\}$ number of classes (randomly chosen from Imagenet dataset) with $nl \in \{3, 5, 7\}$ number of layers. These layers are either convolutional or fully-connected layers only network (FC-only network). For convolutional networks, we iterate between convolutional and max-pooling layers nl times to predict nc classes. For FC-only network, we first flatten the input image, then add nl FC layers. All networks are learned using the RMSProp method (Hinton et al., 2012) with 0.001 (for convolutional) 0.0001 (for FC, smaller rate was critical for learning) learning rate for 100 epochs. The training dataset was prepared with standard data augmentation: feature-wise normalization, linear translation (0.02 range) and horizontal flip.

4.2.2. MORE COMPLEX NETWORK (INCEPTION)

We also measure the closure effect in the more complex and more widely used InceptionV3 network (Szegedy et al., 2016). This network was trained on 1.2 million ImageNet images, with similar augmentation to the simple network: horizontal flips, featurewise normalization, aspect ratio adjustment, shifts, and color distortion. It was trained to top-5 accuracy of 92% over 120 epochs with a batch size of 4096. The learning rate and weights decay followed those of Kornblith et al. (2018).

5. Exploring closure effect: results and discussion

Experimental results relating to each hypothesis are presented, followed by discussion. Our findings are supported by evidence from both simple and complex net-

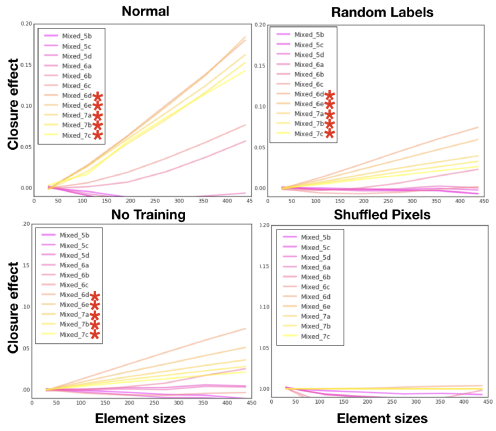

Figure 6. **H1.** Inception Results closely mirror simple network results (Fig. 5)

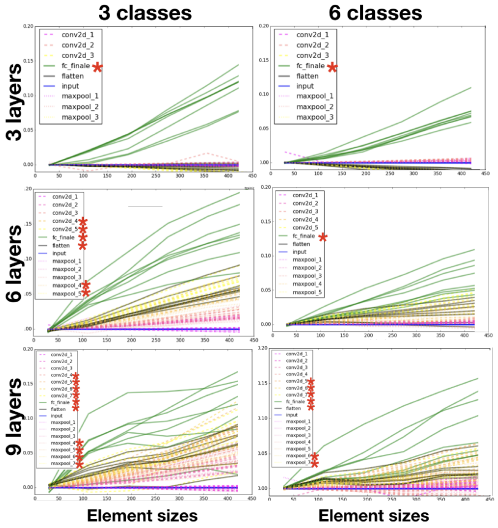

Figure 7. **H2.** Layers close to the prediction layer typically exhibit stronger closure effect. Layers resulted in significance in Welch t-test are marked with $*$.

works (Szegedy et al., 2016).

Reporting significance and multiple runs
Each run produces one significance testing for each element size for each layer. To reduce the clutter, we only report the significance of the largest element size for each layer. When layer l resulted in statistical significance in the difference of closure effect, we mark it with $*$ in the legend. Some plots show multiple runs (8-10, arbitrarily chosen) to ensure consistency. When multiple runs are present, we report the significance results for a randomly chosen run.

5.1. H1. Is the closure effect is associated with generalization?

Results: We observe that the closure effect is influenced by both the network's ability to generalize, as well as its ability to extract some common features. Fig. 5 shows that the NORMAL network has the strongest closure effect, followed by the RANDOM LABELS 1000 network, then RANDOM LABELS nc networks followed by UNTRAINED networks and SMALL DATA network. We also confirm similar patterns in Inception (Fig. 6) and saw the closure effect grow as element size is increased.

Despite some evidence for an association between generalization and a closure effect, readers may wonder about somewhat odd observations: the closure effect in UNTRAINED and the two NNs trained on random labels.

Discussion 1: Why do UNTRAINED networks have the closure effect? The closure effect in UNTRAINED networks aligns with many recent discoveries that suggest that UNTRAINED networks are already a good feature extractor (Ulyanov et al., 2018). When we intentionally decrease features extraction ability (e.g., SMALL DATA) or destroy features that can be extracted by using degenerate data (e.g., SHUFFLED PIXELS or WHITE NOISE), then this sensitivity to the closure effect decreases. We believe this has to do with a network trying to fit the random noise and thus unlearning any useful features it may previously have had (more on this in **H3.**).

Discussion 2: Why does a network trained with random labels have the closure effect? Fig. 5 and Fig. 6 show that even NNs trained with random labels show some closure effect. We conjecture that this is because even though the model cannot truly generalize (as labels are randomly assigned to each image), it can still achieve a very low training error, and in order to do so, it has to use its parameters efficiently. One way to do so is to organize the first layers of the model such that they extract representations that capture what is most commonly seen in the training images, irrespective of their assigned labels. Therefore, we see more closure effect was shown when the network saw more images (RANDOM LABELS 1000) than less (RANDOM LABELS nc). This probably gives rise to the observed closure effect in hidden representations.

Note that it is hard to disentangle a network's ability to extract common representations from its ability to generalize, since we cannot build a network that can generalize without learning good representations. What we observe here is that a network that is able to generalize (e.g., NORMAL) exhibits of the **more** of the closure effect, hinting that the closure effect reflects something beyond simply learning features.

5.2. H2. Is the closure effect stronger in higher layers than lower layers?

Results: Higher layers (close to the prediction layer) in convolutional networks typically show more of the closure

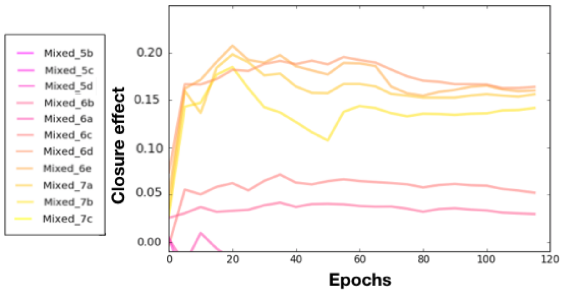

Figure 8. **H3.** The closure effect in the last layer reaches its peak earlier in the training process, then decreases somewhat as it converges. Other layers seem to continue to fluctuate before converging. For degenerated networks, the effect is quickly forgotten. Showing results from Inception.

effect than lower layers (Fig. 6 and Fig. 7). Interestingly, each network seems to have a threshold layer after which all the above layers are statistically significant. In Fig. 7, layers resulting in statistical significance (Welch's t-test, $p < 0.01$) are marked with "*". The more layers a networks has, the more pre-final layers show significance and therefore, the closure effect.

Discussion: This intuitively seems to match with what we know about NN and the key idea of Gestalt psychology: that "the whole is different from the sum of its parts (Koffka, 2013)". There has been some evidence that the lower layers capture lower level features, while higher layers learn higher level features (Bau et al., 2017). If the lower layers extract lower features, in Gestalt terminology, the "parts", then the closure effect must occur more strongly when the "whole" is detected– in the higher layers.

Interestingly, studies using brain recordings in primates also discovered a discrepancy between visual area (V1-V5). von der Heydt et al. (1984) found that neurons that clearly responded to low-contrast figures were silent when illusory stimuli with the same perceived contrast were presented. However, recordings from an area one level higher (V2) of the monkey visual cortex found illusory stimuli responses in about one third of the recorded neurons. They also found response signals get weaker as element size decreases.

5.3. H3. Does closure effect generally increase during training before convergence?

Results: We hypothesized that the closure effect would increase during training and will converge, similar to a typical validation accuracy curve. This varies depending on the layer. In NORMAL network, the closure effect reaches its peak earlier in the iteration, then fluctuates as it learns, then forgets the effect slightly as it converges in both simple network and Inception (Fig. 8). This is typically observed in higher layers (e.g., Mixed 7a and above in Inception). In lower layers (e.g., Mixed 6d and lower), the closure effect increases then converges.

On the other hand, networks trained with degenerate training data (e.g., SHUFFLED PIXELS) start with some closure effect, since UNTRAINED network exhibits some closure effect. However the closure effect drops immediately and stays close to zero in the duration of training.

Discussion: The rapid decrease of the closure effect in networks other than NORMAL aligns with our findings; in the process of trying to fit to random data, the network loses some of the initial feature extraction properties that it had had due to convolutional operations (Ulyanov et al., 2018), and the closure effect is also lost.

The fluctuation during learning before convergence is an interesting symptom, and may benefit from further study. This hints that the closure effect may reflect a prediction-related signal that can be useful (e.g., to determine stopping conditions).

5.4. H4. Is closure effect uninfluenced by simple input manipulations?

In this section, we try to invalidate a hypothesis that a seemingly meaningless input manipulation on the training data (e.g., image brightness) will arbitrarily influence the closure effect. For example, we should not be able to increase/decrease the closure effect by simply making training images brighter/darker (i.e., multiplying/dividing images with constant).

Results: There is no strong pattern between closure effects and brightness of the training images (Fig. 9).

Discussion: Note that the variance in each run reflects the amount of information lost by multiplying or dividing each pixel and saturating them (e.g., multiplying images with 10 cause some images to be no longer identifiable). Naturally in conditions with no strong closure effect (e.g., SHUFFLED PIXELS), there was no difference of the effect from brightness variation.

5.5. H5. Is the closure effect stronger in deeper networks?

The optimal choice of the depth of a network is an unsolved problem, despite many insightful work (Lin et al., 2017; Ba & Caruana, 2014).

Results: As shown in Fig. 9, there is no strong pattern between the closure effect in varying nl for any nc.

Discussion: In hindsight, perhaps **H1** and this hypothesis cannot coexist. Since it is not always true that adding more layers improve the model's performance (Ba & Caruana, 2014), it is reasonable that we cannot influence the closure

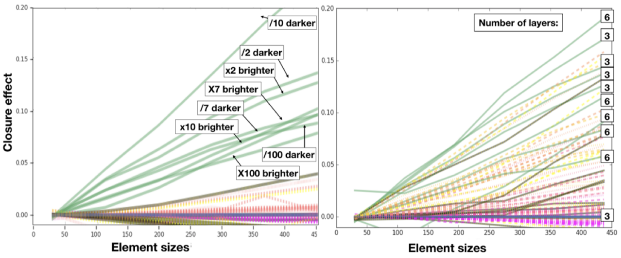

Figure 9. **H4.,H5.** The closure effect does not have clear pattern of change as images become brighter/darker or as the depth of network and/or number of classes trained change. (significance omitted for simplicity)

effect by simply adding more layers.

5.6. H6. Is the closure effect stronger when trained with images intentionally occluded with simple patterns?

Coren & Girgus (1978) noticed that illusory contours generally arise in situations that suggest occlusion. Motivated by this, we trained networks with two types of intentionally occluded training data: 1) randomly chosen number of vertical black bars with random width are drawn in images 2) Bregman B. (Bregman, 2017) type of occlusion added to images (Fig. 2).

Results: Convolutional networks trained with occlusion rarely exhibit more closure effects than normally trained networks. However, in the FC-only case, closure effects do appear to increase when the network is trained with occluded data.

Discussion It may be surprising that training with occlusions does not have a strong effect on our results. One explanation for these results is that the types of occlusions we trained on may not match with the stimuli that we tested on. For example, the stimuli we used does not resemble black bars. In addition, the occlusions we tried in this work are artificial; the Bregman pattern or the vertical bars would rarely occur in real world.

5.7. H7. Is the closure effect stronger with convolutional operators than without?

Results: While both NORMAL convolutional networks and FC-only networks exhibit the closure effect, it is typically stronger in convolutional networks (Fig. 10).

Discussion: Interestingly, in the UNTRAINED condition, we do not observe the closure effect in FC-only networks. As mentioned earlier, this may be due to the fact that UNTRAINED convolutional networks are already good feature extractors (Ulyanov et al., 2018).

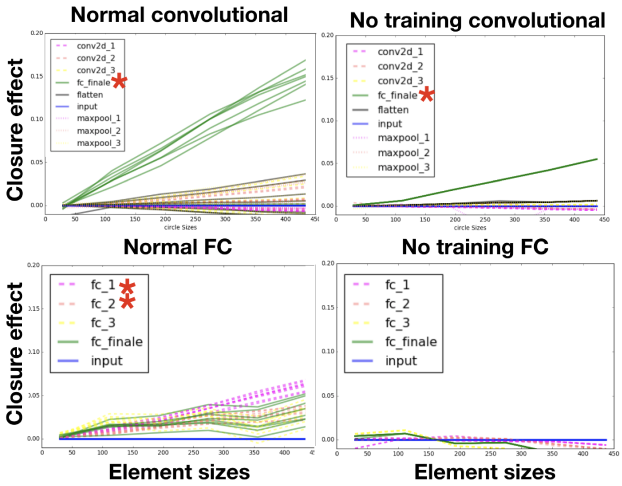

Figure 10. **H7.** Having convolutional layers correlates with stronger closure effect, in UNTRAINED and NORMAL networks. Note that accuracy of FC is much lower than convolutional network. Layers resulted in significance are marked with $*$.

Layer-wise patterns are another interesting thing to notice in this experiment. Unlike in a convolutional network, the last layer (fc_finale, or the green line) does not seem to have the highest closure effect in FC network. Instead, other layers may exhibit the most closure effect (depending on nc and nl). Note that FC-only networks naturally have much lower accuracy than the convolutional networks, which could be a confounding factor.

6. Conclusions

Humans have been studying humans for a long time. The field of psychology has developed useful tools and insights to study human brains– tools that we may be able to borrow to analyze artificial neural networks. In this work, we use one of these tools to study NNs, to gain insights on how similarly or differently they see the world from us. We test for a particular phenomenon, the law of closure, and show that under proper circumstances neural nets also exhibit this effect. The effect seems to be correlated with the network's ability to extract features and to generalize.

The work here is just one step along a much longer path. We believe that exploring other Gestalt laws–and more generally, other psychophysical phenomena–in the context of neural networks is a promising area for future research. Understanding where humans and neural networks differ will be helpful for research on interpretability by enlightening the fundamental differences between the two interesrting species.

References

Ba, J. and Caruana, R. Do deep nets really need to be deep? In *Advances in neural information processing systems*, pp. 2654–2662, 2014.

Barrett, D. G., Hill, F., Santoro, A., Morcos, A. S., and Lillicrap, T. Measuring abstract reasoning in neural networks. *arXiv preprint arXiv:1807.04225*, 2018.

Bau, D., Zhou, B., Khosla, A., Oliva, A., and Torralba, A. Network dissection: Quantifying interpretability of deep visual representations. In *Computer Vision and Pattern Recognition*, 2017.

Behrens, R. R. Art, design and gestalt theory. *Leonardo*, 31 (4):299–303, 1998.

Bender, L. A visual motor gestalt test and its clinical use. *Research Monographs, American Orthopsychiatric Association*, 1938.

Bregman, A. S. Asking the what for question in auditory perception. In *Perceptual organization*, pp. 99–118. Routledge, 2017.

Cadieu, C. F., Hong, H., Yamins, D. L., Pinto, N., Ardila, D., Solomon, E. A., Majaj, N. J., and DiCarlo, J. J. Deep neural networks rival the representation of primate it cortex for core visual object recognition. *PLoS computational biology*, 10(12):e1003963, 2014.

Cichy, R. M., Khosla, A., Pantazis, D., Torralba, A., and Oliva, A. Comparison of deep neural networks to spatio-temporal cortical dynamics of human visual object recognition reveals hierarchical correspondence. *Scientific Reports*, 6:27755, June 2016. doi: 10.1038/srep27755.

Coren, S. and Girgus, J. S. *Seeing is deceiving: The psychology of visual illusions.* Lawrence Erlbaum, 1978.

Desolneux, A., Moisan, L., and Morel, J.-M. *From gestalt theory to image analysis: a probabilistic approach*, volume 34. Springer Science & Business Media, 2007.

Eckstein, M. P., Koehler, K., Welbourne, L. E., and Akbas, E. Humans, but not deep neural networks, often miss giant targets in scenes. *Current Biology*, 27(18):2827–2832, 2017.

Elder, J. and Zucker, S. The effect of contour closure on the rapid discrimination of two-dimensional shapes. *Vision Research*, 33(7):981–991, 1993.

Elsayed, G. F., Shankar, S., Cheung, B., Papernot, N., Kurakin, A., Goodfellow, I., and Sohl-Dickstein, J. Adversarial examples that fool both computer vision and time-limited humans. In *Advances in Neural Information Processing Systems*, 2018.

Erlikhman, G. and Kellman, P. J. Modeling spatiotemporal boundary formation. *Vision research*, 126:131–142, 2016.

Feldman, J. What is a visual object? *Trends in Cognitive Sciences*, 7(6):252–256, 2003.

Geirhos, R., Janssen, D. H., Schütt, H. H., Rauber, J., Bethge, M., and Wichmann, F. A. Comparing deep neural networks against humans: object recognition when the signal gets weaker. *arXiv preprint arXiv:1706.06969*, 2017.

Geisler, W. S. Sequential ideal-observer analysis of visual discriminations. *Psychological review*, 96(2):267, 1989.

Grandchamp, R. and Delorme, A. Single-trial normalization for event-related spectral decomposition reduces sensitivity to noisy trials. *Frontiers in psychology*, 2:236, 2011.

Härdle, W. and Simar, L. *Applied multivariate statistical analysis*, volume 22007. Springer, 2007.

Hinton, G., Srivastava, N., and Swersky, K. Neural networks for machine learning lecture 6a overview of mini-batch gradient descent. *Cited on*, pp. 14, 2012.

Jkel, F., Singh, M., Wichmann, F. A., and Herzog, M. H. An overview of quantitative approaches in gestalt perception. *Vision Research*, 126:3 – 8, 2016. ISSN 0042-6989. doi: https://doi.org/10.1016/j.visres.2016.06.004. URL http://www.sciencedirect.com/science/article/pii/S0042698916300475. Quantitative Approaches in Gestalt Perception.

Keil, A., Müller, M. M., Ray, W. J., Gruber, T., and Elbert, T. Human gamma band activity and perception of a gestalt. *Journal of Neuroscience*, 19(16):7152–7161, 1999.

Keren, A. S., Yuval-Greenberg, S., and Deouell, L. Y. Saccadic spike potentials in gamma-band eeg: characterization, detection and suppression. *Neuroimage*, 49(3): 2248–2263, 2010.

Kimchi, R. Primacy of wholistic processing and global/local paradigm: a critical review. *Psychological bulletin*, 112 (1):24, 1992.

Kimchi, R., Yeshurun, Y., Spehar, B., and Pirkner, Y. Perceptual organization, visual attention, and objecthood. *Vision Research*, 126:34 – 51, 2016. ISSN 0042-6989. doi: https://doi.org/10.1016/j.visres.2015.07.008. URL http://www.sciencedirect.com/science/article/pii/S0042698915003119. Quantitative Approaches in Gestalt Perception.

Koffka, K. *Principles of Gestalt psychology*. Routledge, 2013.

Kornblith, S., Shlens, J., and Le, Q. V. Do better imagenet models transfer better? *CoRR*, abs/1805.08974, 2018. URL http://arxiv.org/abs/1805.08974.

Kriegeskorte, N., Mur, M., and Bandettini, P. A. Representational similarity analysis-connecting the branches of systems neuroscience. *Frontiers in systems neuroscience*, 2:4, 2008.

Lin, H. W., Tegmark, M., and Rolnick, D. Why does deep and cheap learning work so well? *Journal of Statistical Physics*, 168(6):1223–1247, 2017.

Luck, S. J. *An introduction to the event-related potential technique*. MIT press, 2014.

Ma, W. J. Organizing probabilistic models of perception. *Trends in cognitive sciences*, 16(10):511–518, 2012.

Machilsen, B., Wagemans, J., and Demeyer, M. Quantifying density cues in grouping displays. *Vision research*, 126: 207–219, 2016.

Maess, B., Schröger, E., and Widmann, A. High-pass filters and baseline correction in m/eeg analysis. commentary on:how inappropriate high-pass filters can produce artefacts and incorrect conclusions in erp studies of language and cognition. *Journal of neuroscience methods*, 266: 164–165, 2016.

Plöchl, M., Ossandón, J. P., and König, P. Combining eeg and eye tracking: identification, characterization, and correction of eye movement artifacts in electroencephalographic data. *Frontiers in human neuroscience*, 6:278, 2012.

Rajalingham, R., Issa, E. B., Bashivan, P., Kar, K., Schmidt, K., and DiCarlo, J. J. Large-scale, high-resolution comparison of the core visual object recognition behavior of humans, monkeys, and state-of-the-art deep artificial neural networks. *Journal of Neuroscience*, 38(33):7255–7269, 2018.

Ren, X. and Malik, J. Learning a classification model for segmentation. In *null*, pp. 10. IEEE, 2003.

Riesenhuber, M. and Poggio, T. Hierarchical models of object recognition in cortex. *Nature neuroscience*, 2(11): 1019, 1999.

Ringach, D. L. and Shapley, R. Spatial and temporal properties of illusory contours and amodal boundary completion. *Vision research*, 36(19):3037–3050, 1996.

Ritter, S., Barrett, D. G., Santoro, A., and Botvinick, M. M. Cognitive psychology for deep neural networks: A shape bias case study. *arXiv preprint arXiv:1706.08606*, 2017.

Sanguinetti, J. L., Trujillo, L. T., Schnyer, D. M., Allen, J. J., and Peterson, M. A. Increased alpha band activity indexes inhibitory competition across a border during figure assignment. *Vision Research*, 126:120–130, 2016.

Schultz, D. P. and Schultz, S. E. *A history of modern psychology*. Cengage Learning, 2015.

Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., and Wojna, Z. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2818–2826, 2016.

Tanner, D., Norton, J., Morgan-Short, K., and Luck, S. J. On high-pass filter artifacts (theyre real) and baseline correction (itsa good idea) in erp/ermf analysis. *Journal of neuroscience methods*, 266:166–170, 2016.

Ulyanov, D., Vedaldi, A., and Lempitsky, V. Deep image prior. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 9446–9454, 2018.

von der Heydt, R., Peterhans, E., and Baumgartner, G. Illusory contours and cortical neuron responses. *Science*, 224(4654):1260–1262, 1984.

Wagemans, J., Elder, J. H., Kubovy, M., Palmer, S. E., Peterson, M. A., Singh, M., and von der Heydt, R. A century of gestalt psychology in visual perception: I. perceptual grouping and figure–ground organization. *Psychological bulletin*, 138(6):1172, 2012a.

Wagemans, J., Feldman, J., Gepshtein, S., Kimchi, R., Pomerantz, J. R., van der Helm, P. A., and van Leeuwen, C. A century of gestalt psychology in visual perception: Ii. conceptual and theoretical foundations. *Psychological bulletin*, 138(6):1218, 2012b.

Welch, B. L. The generalization of student's' problem when several different population variances are involved. *Biometrika*, 34(1/2):28–35, 1947.

Wertheimer, M. Laws of organization in perceptual forms. *A source book of Gestalt Psychology*, 1923.

Westheimer, G. Gestalt theory reconfigured: Max wertheimer's anticipation of recent developments in visual neuroscience. *Perception*, 28(1):5–15, 1999.

Yamins, D. L., Hong, H., Cadieu, C. F., Solomon, E. A., Seibert, D., and DiCarlo, J. J. Performance-optimized hierarchical models predict neural responses in higher visual cortex. *Proceedings of the National Academy of Sciences*, 111(23):8619–8624, 2014.

Zinker, J. *Creative process in Gestalt therapy*. Brunner/Mazel, 1977.