

ARE THERE ANY ‘OBJECT DETECTORS’ IN THE HIDDEN LAYERS OF CNNs TRAINED TO IDENTIFY OBJECTS OR SCENES?

Anonymous authors

Paper under double-blind review

ABSTRACT

Various methods of measuring unit selectivity have been developed with the aim of better understanding how neural networks work. But the different measures provide divergent estimates of selectivity, and this has led to different conclusions regarding the conditions in which selective object representations are learned and the functional relevance of these representations. In an attempt to better characterize object selectivity, we undertake a comparison of various selectivity measures on a large set of units in AlexNet, including localist selectivity (Bowers et al., 2014), precision (Zhou et al., 2015), class-conditional mean activity selectivity (CCMAS) (Morcos et al., 2018), network dissection (Zhou et al., 2018a), the human interpretation of activation maximization (AM) images, and standard signal-detection measures. We find that the different measures provide different estimates of object selectivity, with precision and CCMAS measures providing misleadingly high estimates. Indeed, the most selective units had a poor hit-rate or a high false-alarm rate (or both) in object classification, making them poor object detectors. We fail to find any units that are even remotely as selective as the ‘grandmother cell’ units reported in recurrent neural networks. In order to generalize these results, we compared selectivity measures on a few units in VGG-16 and GoogLeNet trained on the ImageNet or Places-365 datasets that have been described as ‘object detectors’. Again, we find poor hit-rates and high false-alarm rates for object classification.

1 INTRODUCTION

There have been recent attempts to understand how neural networks (NNs) work by analyzing hidden units one-at-a-time using various measures such as localist selectivity (Bowers et al., 2014), class-conditional mean activity selectivity (CCMAS) (Morcos et al., 2018), precision (Zhou et al., 2015), network dissection (Zhou et al., 2018a), and activation maximization (AM) (Erhan et al., 2009b). These measures are all taken to provide evidence that some units respond highly selectively to categories of objects under some conditions. Not only are these findings surprising given the widespread assumption that NNs learned highly distributed and entangled representations, they raise a host of questions, including the functional importance of these selective representations (Zhou et al., 2018b), the conditions in which they are learned (e.g., Morcos et al., 2018), and the relation between these representations and the selective neurons observed in cortex (Bowers, 2009).

To answer these question, it is necessary to have a better understanding of what these metrics actually measure, and how they relate to one another. Accordingly, we directly compare these measures of selectivity on the same set of units as well as adopt standard signal-detection measures in an attempt to provide better measures of single-unit selectivity to object category. In addition, to provide a more intuitive assessment of selectivity, we report jitter-plots for a few of the most selective units that visually display how the unit responds to the different image categories. We focus on AlexNet (Krizhevsky et al., 2012) trained on ImageNet (Deng et al., 2009) because many authors have studied the selectivity of single hidden units in this model using a range of quantitative (Zhou et al., 2018a; 2015) and qualitative (Nguyen et al., 2017; Yosinski et al., 2015; Simonyan et al., 2013) methods. But we also compare different selectivity measures on specific units in VGG-16 (Simonyan and Zisserman, 2014) and GoogLeNet (Szegedy et al., 2015) trained on the the ImageNet and Places-365

datasets that were characterized by Zhou et al. (2018a) as “object detectors” based on their Network Dissection method (Zhou et al., 2018a). Our main findings are:

1. The precision and CCMAS measures are misleading with near-maximum selectivity scores associated with units that strongly respond to many different image categories. By contrast, the signal-detection measures more closely capture the level of selectivity displayed in the jitter-plots (Sec. 3.1).
2. Units with interpretable AM images do not correspond to highly selective representations (Sec. 3.2).
3. The Network Dissection method also provides a misleading measure for “object detectors” (Sec. 3.3).

In one line of research, Bowers et al. (2014; 2016) assessed the selectivity of single hidden units in recurrent neural networks (RNNs) designed to model human short-term memory. They reported many ‘localist’ or ‘grandmother cell’ units that were 100% selective for specific letters or words, where all members of the selective category were more active than and disjoint from all non-members, as can be shown in jitterplots (Berkeley et al., 1995) (see Fig. 1 for a unit selective to the letter ‘j’). The authors argued that the network learned these representations in order to co-activate multiple letters or words at the same time in short-term memory without producing ambiguous blends of overlapping distributed patterns (the so-called ‘superposition catastrophe’). Consistent with this hypothesis, localist units did not emerge when the model was trained on letters or words one-at-a-time (Bowers et al., 2014) (see Fig. 1 for an example of a non-selective unit).

In parallel, researchers have reported selective units in the hidden layers of various CNNs trained to classify images into one of multiple categories (Zhou et al., 2015; Morcos et al., 2018; Zeiler and Fergus, 2014; Erhan et al., 2009b), for a review see Bowers (2017). For example, Zhou et al. (2015) assessed the selectivity of units in the pool5 layer of two CNNs trained to classify images into 1000 objects and 205 scene categories, respectively. They reported many highly selective units that they characterized as ‘object detectors’ in both networks. Similarly, Morcos et al. (2018) reported that CNNs trained on CIFAR-10 and ImageNet learned many highly selective hidden units, with CCMAS scores often approaching the maximum of 1.0. These later findings appear to be inconsistent with Bowers et al. (2016) who failed to observe selective representations in fully connected NNs trained on stimuli one-at-a-time (see Fig. 1), but the measures of selectivity that have been applied across studies are different, and accordingly, it is difficult to directly compare results.

A better understanding of the relation between selectivity measures is vital given that different measures are frequently used to address similar issues. For example, both the human interpretability of generated images (Le, 2013) and localist selectivity (Bowers et al., 2014) have been used to make claims about ‘grandmother cells’, but it is not clear whether they provide similar insights into unit selectivity. Similarly, based on their precision metric, Zhou et al. (2015) claim that the object detectors learned in CNNs play an important role in identifying specific objects, whereas Morcos et al. (2018) challenge this conclusion based on their finding that units with high CCMAS measures were not especially important in the performance of their CNNs and concluded: “...it implies that methods for understanding neural networks based on analyzing highly selective single units, or finding optimal inputs for single units, such as activation maximization (Erhan et al., 2009a) may be misleading”. This makes a direct comparison between selectivity measures all the more important.

In order to directly compare and have a better understanding of the different selectivity measures we assessed (1) localist, (2) precision, and (3) CCMAS selectivity of the conv5, fc6, and fc7 of AlexNet trained on ImageNet, and in addition, we employed a range of signal detection methods on these units, namely, (4) recall with 100% and 95% precision, (5) maximum informedness, (6) specificity at maximum informedness, and (7) recall (also called *sensitivity*) at maximum informedness, and false alarm rates at maximum informedness (described in Sec. 2). We also assessed the selectivity of a few units in VGG-16 and GoogLeNet models trained on the ImageNet and Places-365 dataset that were highly selective according to the Network Dissection method (Zhou et al., 2018a). We show that the precision and CCMAS measures often provide misleadingly high estimates of object selectivity compared to other measures, and we do not find any units that can be reasonably described as ‘object detectors’ given that the most selective units show a low hit-rate or a high false-alarm rate (or both) when classifying images. At best, the most selective units in CNNs are sensitive to some unknown feature that is frequently (but not exclusively) associated with the class in question.

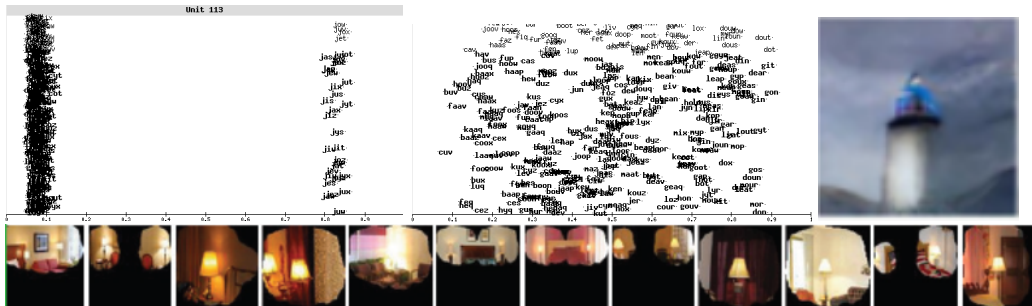


Figure 1: Examples of selectivity measures used. **Top left:** jitterplot of unit 113 in an RNN (under the superposition constraint) selective to the letter ‘j’ (Bowers et al., 2016). **Top middle:** jitterplot of a non-selective unit 160 found in an RNN trained on words one-at-a-time from (Bowers et al., 2016). **Top right:** Activation maximization image of unit conv5₀ AlexNet that resembles a lighthouse (Nguyen et al., 2016). **Bottom:** highest-activation images for a ‘lamp’ detector with 84% precision in the layer conv5 of AlexNet; from (Zhou et al., 2015).

In addition to these quantitative measures and jitter-plots we assessed a selectivity with a common qualitative measure, namely, human interpretation of images generated by a state-of-the-art activation maximization (AM) method (Nguyen et al., 2017). AM images are generated to strongly activate individual units, and some of them are interpretable by humans (e.g., a generated image that looks like a lighthouse, see Fig. 1). For the first time, we systematically evaluated the interpretability of the AM images and compare these ratings with the selectivity measures for corresponding units. We show that the few hidden units with interpretable AM images are not highly selective.

2 METHODS

Network and Dataset All ~ 1.3 M photos from the ImageNet ILSVRC 2012 dataset (Deng et al., 2009) were cropped to 277×277 pixels and classified by the pre-trained AlexNet CNN (Krizhevsky et al., 2012) shipped with Caffe (Jia et al., 2014), resulting in 721,536 correctly classified images. Once classified, the images are not re-cropped nor subject to any changes. We analyzed the fully connected (fc) layers: fc6 and fc7 (4096 units), and the top convolutional layer conv5 which has 256 filters. We only recorded the activations of correctly classified images. The activation files are stored in .h5 format and will be available at <http://anonymizedForReview>. We selected 233 conv5, 2738 fc6, 2239 fc7 units for analysis.

Localist selectivity Following Bowers et al. (2014), we define a unit to be localist for class A if the set of activations for class A was higher and disjoint with those of $\neg A$. Localist selectivity is easily depicted with jitterplots (Berkeley et al., 1995) in which a scatter plot for each unit is generated (see Figs. 1 and 3). Each point in a plot corresponds to a unit’s activation in response to a single image, and only correctly classified images are plotted. The level of activation is coded along the x -axis, and an arbitrary value is assigned to each point on the y -axis.

Precision Precision refers to the proportion of items above some threshold from a given class. The precision method of finding object detectors involves identifying a small subset of images that most strongly activate a unit and then identifying the critical part of these images that are responsible for driving the unit. Zhou et al. (2015) took the 60 images that activated a unit the most strongly and asked independent raters to interpret the critical image patches (e.g., if 50 of the 60 images were labeled as ‘lamp’, the unit would have a precision index of 50/60 or 83%; see Fig. 1). Object detectors were defined as units with a precision score $> 75\%$: they reported multiple such detectors. Here, we approximate this approach by considering the 60 images that most strongly activate a given unit and assess the highest percentage of images from a given output class.

CCMAS Morcos et al. (2018) introduced a selectivity index called the Class-conditional Mean Activation Selectivity (CCMAS). The CCMAS for class A compares the mean activation of all images in class A , μ_A , with the mean activation of all images not in class A , $\mu_{\neg A}$, and is given by: $(\mu_A - \mu_{\neg A}) / (\mu_A + \mu_{\neg A})$. Here, we assessed class selectivity for the highest mean activation class.

Activation Maximization We harnessed an activation maximization method called Plug & Play Generative Networks (Nguyen et al., 2017) in which an image generator network was used to generate images (AM images) that highly activate a unit in a target network. We used the public code released by Nguyen et al. (2017) and their default hyperparameters.¹ We generated 100 separate images that maximally activated each unit in the conv5, fc6, and fc8 layers of AlexNet and asked participants to judge whether they could identify any repeating objects, animals, or places in images in a behavioral experiment (Sec. 3.2). Readers can test themselves at: <https://research.sc/participant/login/dynamic/63907FB2-3CB9-45A9-B4AC-EFFD4C4A95D5>

Recall with perfect and 95% precision Recall with perfect and 95% precision are related to localist selectivity except that they provide a continuous rather than discrete measure. For recall with perfect precision we identified the image that activated a given unit the most and counted the number of images from the same class that were more active than all images from all other classes. We then divided this result by the total number of correctly identified images from this class. A recall with a perfect precision score of 1 is equivalent to a localist representation. Recall with a 95% precision allows 5% false alarms.

Maximum informedness Maximum informedness identifies the class and threshold where the highest proportion of images above the threshold and the lowest proportion of images below the threshold are from that class (Powers, 2003). The informedness is computed for each class at each threshold, with the highest value selected. Informedness summarises the diagnostic performance of unit for a given class at a certain threshold based on the recall [True Positives / (True Positives + False Negatives)] and specificity [True Negatives / (True Negatives + False Positives)] in the formula [informedness = recall + specificity - 1] (Powers, 2003).

Sensitivity or Recall at Maximum Informedness For the threshold and class selected by Maximum Informedness, recall (or hit-rate) is the proportion of items from the given class that are above the threshold. Also known as true positive rate.

Specificity at Maximum Informedness For the threshold and class selected by Maximum Informedness, the proportion of items that are not from the given class that are below the threshold. Also known as true negative rate.

False Alarm Rate at Maximum Informedness For the threshold and class selected by Maximum Informedness, the proportion of items that are not from the given class that are above the threshold.

Network Dissection To assess the selectivity of a unit in the Network Dissection technique, Zhou et al. (2018a) compute the Intersection over Union (IoU) of an annotated input image L_c , for the set of all ‘concepts’ c and a spatial activation map, M_k , of where a unit k is. A unit k is taken as a detector for concept c if its $\text{IoU}_{k,c}$ exceeds a pre-defined threshold T . See Zhou et al. (2018a) for more details.

3 RESULTS

3.1 COMPARISON OF SELECTIVITY MEASURES IN ALEXNET

The results from the various of selectivity measures applied to the conv5, fc6, and fc7 layers of AlexNet are displayed in Fig. 2a–i. We did not plot the localist selectivity as there were no localist ‘grandmother units’. The first point to note is that multiple units in the fc6 and fc7 layers had near 100% precision scores and multiple units had CCMAS scores approaching 1. For example, in layer fc7, we found 14 units with a precision > 0.9 , and 1487 units with a CCMAS > 0.9 . The second point is that other measures provided much reduced estimates of selectivity. For example, the unit with the highest recall with a perfect precision score was only .08 (unit 255 responding to images of Monarch butterflies), and the unit with the top maximum informedness score (unit 3290 also responding to images of Monarch butterflies with a score of 0.91) had a false alarm rate above its optimal threshold $> 99\%$ (indeed the minimum false alarm rate was 0.96).

To illustrate the contrasting measures of selectivity consider unit fc6₁₁₉₉ depicted in Fig. 3 that has a precision score of 98% and a CCMAS score of .92. By Zhou et al.’s criterion, this is a ‘Monarch Butterfly’ detector (its precision score $> 75\%$). By contrast, the scatter plot and signal-detection

¹<https://github.com/Evolving-AI-Lab/ppgn>

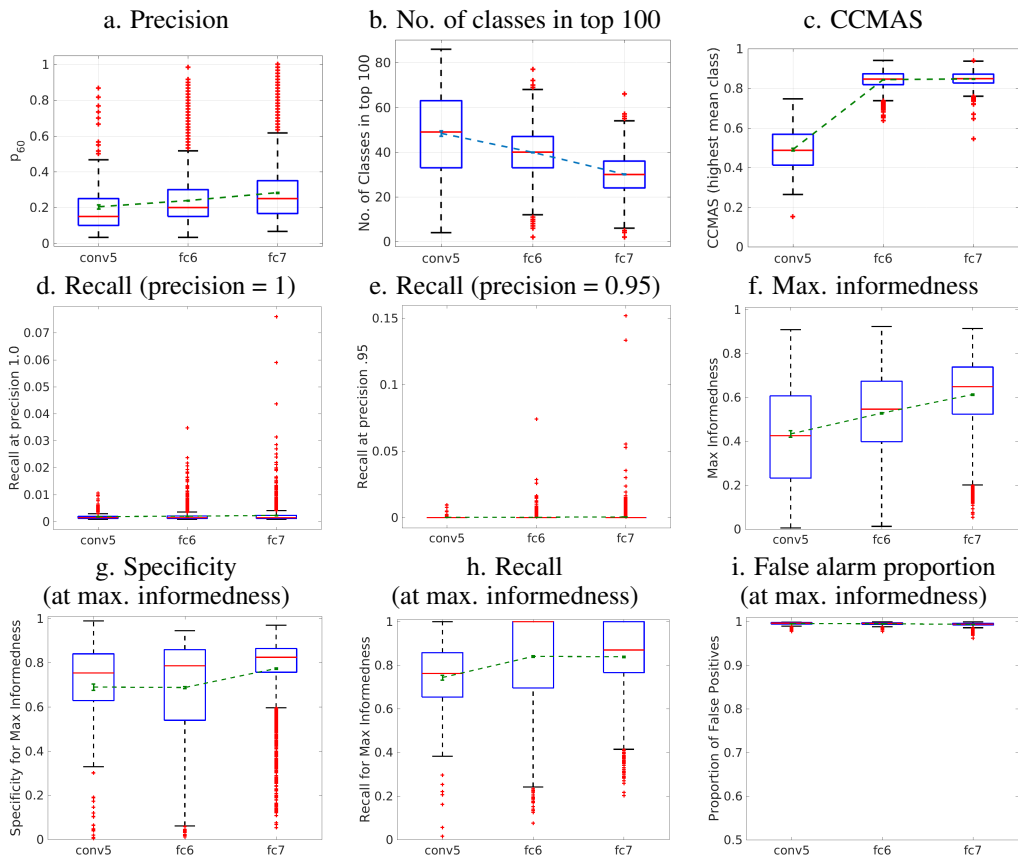


Figure 2: Different selectivity measures across the conv5, fc6, and fc7 layers of AlexNet. Red-line: median of data, top and bottom of box edges is the 25th and 75th percentile, whiskers extend to extreme edges of distribution not considered outliers and red crosses are outliers. Green points and dashed lines are the means of the distributions with standard errors. The high levels of selectivity observed with the precision and CCMAS measures are in stark contrast with the low levels of selectivity observed with the recall with perfect precision and high false-alarm rates at maximum informedness.

scores show this is a mischaracterisation of this unit given that the false alarm rate at maximum informedness was greater than 99% and the modal response to Monarch butterflies was zero.

What level of selectivity is required before a unit can be considered an ‘object detector’ for a given category? In the end, this is a terminological point. On an extreme view, one might limit the term to the ‘grandmother units’ that categorize objects with perfect recall and specificity. But at minimum, it would only seem reasonable to describe a unit as a detector for a specific object category if there is some threshold of activation that supports more hits than misses (the unit is strongly activated by the majority of images from a given category), and at the same time, supports more hits than false alarms (the unit is strongly activated by items from the given category more often than by items from other categories). A suitable selectivity measure should reflect any trade-off between these two. The signal-detection measures show that the most selective units do not even come close to this low bar. That said, some units respond strongly to some (unknown) features that are weakly correlated with an object category. For instance, unit fc6₁₁₉₉ is responding to features that occur more frequently in Monarch Butterflies than other categories. This can also be seen in a recent ablation study in which removing the most selective units tended to impair the CNN’s performance in identifying the corresponding object categories more than other categories (Zhou et al., 2018b). But again, the pattern of performance was not consistent with the units being labeled ‘object detectors’.




Figure 3: Data for unit $fc6_{1199}$. **Left:** activation jitterplot, black diamonds: Monarch butterfly images; grey circles: all other classes; white dashed line: threshold for the butterfly class maximum informedness; blue solid line: threshold for top 60 activations. **Middle:** histogram of activations of Monarch butterflies. **Inset:** zoomed-in histogram of all activations across all ImageNet classes of unit $fc6_{1199}$ (N.B. this plot shows only the highest 121,586 activations; there are 596,734 activations at 0). There are Monarch butterfly images covering the whole range of values, with 72 images (5.8% of the total) having an activation of 0. **Right:** example ImageNet images with activations of 0 (top), the mean, 39.2 ± 0.6 , (middle), and the maximum, 95, (bottom) of the range. Although the high precision score suggests that this unit is a butterfly detector this is misleading given there are butterfly images over the entire activation range (including 0).

3.2 HUMAN INTERPRETATION OF ACTIVATION MAXIMIZATION IMAGES FOR ALEXNET UNITS

Activation Maximization is one of the most commonly used interpretability methods for explaining what a single unit has learned in many artificial CNNs and even biological neural networks (see Nguyen et al. (2019) for a survey). Our behavioral experiment provides the first quantitative assessment of AM images and compares AM interpretability to other selectivity measures.

Table 1: Human judgements of whether AM images look like familiar objects in layers conv5, fc6, and fc8 in AlexNet.

layer	% ‘yes’ responses (a)	% units $\geq 80\%$ ‘yes’ response (b)	% overlap between humans and:		
			humans (c)	most active object (d)	CCMAS class (e)
conv5	21.7% ($\pm 1.1\%$)	4.3% ($\pm 1.3\%$)	89.5% ($\pm 5.7\%$)	34.1% ($\pm 14.4\%$)	0%
fc6	21.0% ($\pm 0.4\%$)	3.1% ($\pm 0.4\%$)	80.4% ($\pm 4.1\%$)	23.3% ($\pm 5.9\%$)	18.9% ($\pm 5.9\%$)
fc8	71.2% ($\pm 0.6\%$)	59.3% ($\pm 1.6\%$)	96.5% ($\pm 0.4\%$)	95.4% ($\pm 0.6\%$)	94.6% ($\pm 0.7\%$)




Figure 4: Example AM images that were either judged by all participants to contain objects (a–c) or to be uninterpretable as objects (d–f). The human label for $conv5_{183}$ (a) was ‘dogs’; the most active image was of a ‘flat-coated retriever’; CCMAS class was ‘monitor’. For $fc6_{319}$ (b), subjects reported ‘green peppers’ or ‘apples’ (all classified as the same broad class in our analysis); both the most active item and CCMAS class were ‘Granny Smith apples’. For $fc8_{969}$ (c), humans suggested ‘beverage’ or ‘drink’; both the most active item and CCMAS class were ‘eggnog’.

We generated 100 AM images for every unit in the layers conv5, fc6, and fc8 in AlexNet, as in Nguyen et al. (2017), and displayed them as 10×10 -image panels. A total of 3,299 image panels were used in the experiment (995 fc8, 256 conv5, and 2048 randomly selected fc6 image panels) and were divided into 64 counterbalanced lists for testing. To assess the interpretability for these units as object detectors, 333 paid volunteers were asked to look at image panels and asked if the images had an object / animal or place in common. If the answer was ‘yes’, they were asked to write down a generic name for that object (e.g. “fish” rather than “goldfish”). Analyses of common responses was done for any units where over 80% of humans agreed there was an object present.

The results are summarized in Table 1. Not surprisingly, the AM images for output fc8 units are the most human-recognizable as objects across the AlexNet layers (71.2%; Table 1a). In addition, when they were given a consistent interpretation, they almost always (95.4%; Table 1d) match the corresponding ImageNet category. By contrast, less than 5% of units in conv5 or fc6 were associated with consistently interpretable images (Table 1b), and the interpretations only weakly matched the category associated with the highest-activation images or CCMAS selectivity (Table 1d–e). Apart from showing that there are few interpretable units in the hidden layers of AlexNet, our findings show that the interpretability of images does not imply a high level of selectivity given the signal-detection results (Fig. 2d–h). See Fig. 4 for an example of the types of images that participants rated as objects or non-objects.

3.3 COMPARING SELECTIVITY MEASURES IN OTHER CNNs

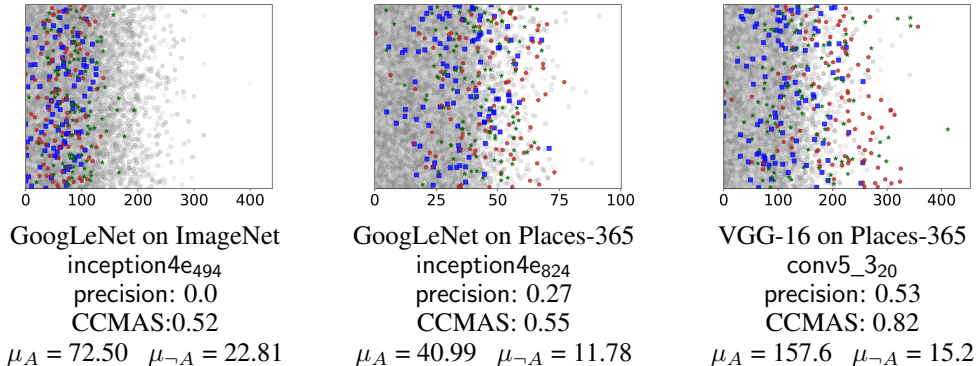


Figure 5: The units with with the highest Network Dissection scores for the category ‘bus’. The scatter plots, precision, and CCMAS scores all indicate a low selectivity for this category. **blue** squares: ‘school bus’; **red** pentagons: ‘trolleybus’; **green** stars: ‘minibus’; **grey** circles: other classes.

Thus far we have assessed the selectivity of hidden units in AlexNet and shown that no units can reasonably be characterized as object detectors despite the high precision and CCMAS scores of some units. This raises the question as to whether more recent CNNs learn object detector units. In order to address this we display jitterplots for three units that have the highest IoU scores according to the Network Dissection for the category BUS in GoogLeNet trained on ImageNet, GoogLeNet trained on Places-365, and VGG-16 trained on Places-365, respectively (Zhou et al., 2018a). Models trained on the Places-365 dataset learn to categorize images into scenes (e.g., bedrooms, kitchens, etc.) rather than into object categories, and nevertheless, Zhou et al. (2018a) reported more object detectors in the former models. We illustrate the selectivity of the BUS category because it is an output category in ImageNet so we can easily plot the jitterplots for these units.

As was the case with AlexNet, the jitterplots show that the most selective units show some degree of selectivity, with the BUS images more active on average compared to non-Buses, and the percentage of nonzero activations for BUS higher than the non-BUS categories (see tables 4 - 6 in the appendix for summary of more units). But the units are no more selective than the units we observed in AlexNet. Indeed, the precision measure of selectivity for the first units is 0.0, with none of the units having a precision of .75 that was the criterion of object detectors by Zhou et al. (2015), and CCMAS scores for first two units were roughly similar to the mean CCMAS score for AlexNet units in conv 5 (and much lower than the mean in fc6 and fc7). The most selective VGG-16 unit trained on Places-365 has

lower precision and CCMAS scores than the Monarch Butterfly unit depicted in Figure 3. So again, different measures of selectivity provide support different conclusions, and even the most selective units are far from the selective units observed in recurrent networks as reported in Figure 1a. See tables 4 - 6 in the appendix for more details about these three units.

4 DISCUSSIONS AND CONCLUSIONS

Our central finding is that different measures of single-unit selectivity for objects support very different conclusions when applied to the same units in AlexNet. In contrast with the precision (Zhou et al., 2015) and CCMAS (Morcos et al., 2018) measures that suggest some highly selective units for objects in layers conv5, fc6, and fc7, the recall with perfect precision and false alarm rates at maximum informedness show low levels of selectivity. Indeed, the most selective units have a poor hit-rate or a high false-alarm rate (or both) for identifying an object class. The same outcome was observed with units in VGG-16 and GoogLeNet trained on either ImageNet or the Places-365 dataset.

Not only do the different measures provide very different assessments of selectivity, the precision, CCMAS, and Network Dissection measures provide highly misleading estimates of selectivity that have led to mistaken conclusions. For example, unit fc6₁₁₉₉ in AlexNet trained on ImageNet is considered an Monarch Butterfly detector according to Zhou et al. (2015) with a precision score of 98% (and a CCMAS score of .93). But the jitterplot in Fig. 3 and signal detection scores (e.g., high false alarm rate at maximum informedness) show this is a mischaracterisation of this unit. In the same way, the Network Dissection method identified many object detectors in VGG-16 and GoogLeNet CNNs, but the jitterplots in Fig. 5 (and precision scores) show that this conclusion is unjustified. Similarly, the images generated by Activation Maximization also provided a misleading estimate of selectivity given that interpretable images were associated with very low selectivity scores. We would like to emphasize that the results by Morcos et al. (2018) and Zhou et al. (2015; 2018a) are important because they have shown there is some interpretability associated with the units. But our signal-detection results and jitterplots highlight the extent to which objects are encoded in a distributed rather than disentangled manner in CNNs.

What should be made of the finding that localist representations are sometimes learned in RNNs (units with perfect specificity and recall), but not in AlexNet and related CNNs? The failure to observe localist units in the hidden layers of these CNNs is consistent with Bowers et al. (2014)’s claim that these units emerge in order to support the co-activation of multiple items at the same time in short-term memory. That is, localist representations may be the solution to the superposition catastrophe, and these CNNs only have to identify one image at a time. The pressure to learn highly selective representations in response to the superposition constraint may help explain the reports of highly selective neurons in cortex given that the cortex needs to co-activate multiple items at the same time in order to support short-term memory (Bowers et al., 2016).

Note, the RNNs that learned localist units were very small in scale compared to CNNs we have studied here, and accordingly, it is possible that the contrasting results reflect the size of the networks rather than the superposition catastrophe *per se*. Relevant to this issue a number of authors have reported the existence of selective units in larger RNNs with long-short term memory (LSTM) units (Karpathy et al., 2016; Radford et al., 2017; Lakretz et al., 2019; Na et al., 2019). Indeed, Lakretz et al. (2019) use the term ‘grandmother cell’ to describe the units they observed. It will be interesting to apply our measures of selectivity to these larger RNNs and see whether these units are indeed ‘grandmother units’.

It should also be noted that there are recent reports of impressively selective representations in Generative Adversarial Networks (Bau et al., 2019) and Variational Autoencoders (Burgess et al., 2018) where the superposition catastrophe is not an issue. Again, it will be interesting to assess the selectivity of these units according to signal detection measures in order to see whether there are additional computational pressures to learn highly selective or even grandmother cells. We will be exploring these issues in future work. Whatever the case, the main finding here is that CNNs trained to identify objects and scenes one-at-a-time did not learn highly selective object detectors.

REFERENCES

- D. Bau, J.-Y. Zhu, H. Strobelt, B. Zhou, J. B. Tenenbaum, W. T. Freeman, and A. Torralba. Visualizing and understanding generative adversarial networks. In *International Conference on Learning Representations*, 2019. URL https://openreview.net/forum?id=Hyg_X2C5FX.
- I. S. Berkeley, M. R. Dawson, D. A. Medler, D. P. Schopfloch, and L. Hornsby. Density plots of hidden value unit activations reveal interpretable bands. *Connection Science*, 7(2):167–187, 1995.
- J. S. Bowers. On the biological plausibility of grandmother cells: implications for neural network theories in psychology and neuroscience. *Psychological review*, 116(1):220, 2009.
- J. S. Bowers. Grandmother cells and localist representations: a review of current thinking. *Language, Cognition, and Neuroscience*, pages 257–273, 2017.
- J. S. Bowers, I. I. Vankov, M. F. Damian, and C. J. Davis. Neural networks learn highly selective representations in order to overcome the superposition catastrophe. *Psychological review*, 121(2): 248–261, 2014.
- J. S. Bowers, I. I. Vankov, M. F. Damian, and C. J. Davis. Why do some neurons in cortex respond to information in a selective manner? insights from artificial neural networks. *Cognition*, 148:47–63, 2016.
- C. P. Burgess, I. Higgins, A. Pal, L. Matthey, N. Watters, G. Desjardins, and A. Lerchner. Understanding disentangling in β -VAE. *arXiv preprint arXiv:1804.03599*, 2018.
- J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pages 248–255. IEEE, 2009.
- D. Erhan, Y. Bengio, A. Courville, and P. Vincent. Visualizing higher-layer features of a deep network. *University of Montreal*, 1341(3):1, 2009a.
- D. Erhan, Y. Bengio, A. Courville, and P. Vincent. Visualizing higher-layer features of a deep network. *University of Montreal*, 1341(3):1, 2009b.
- Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional architecture for fast feature embedding. *arXiv preprint arXiv:1408.5093*, 2014.
- A. Karpathy, J. Johnson, and L. Fei-Fei. Visualizing and understanding recurrent networks. *Workshop Track at International Conference on Learning Representations*, 2016.
- A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- Y. Lakretz, G. Kruszewski, T. Desbordes, D. Hupkes, S. Dehaene, and M. Baroni. The emergence of number and syntax units in lstm language models. *arXiv preprint arXiv:1903.07435*, 2019.
- Q. V. Le. Building high-level features using large scale unsupervised learning. In *2013 IEEE international conference on acoustics, speech and signal processing*, pages 8595–8598. IEEE, 2013.
- A. S. Morcos, D. G. Barrett, N. C. Rabinowitz, and M. Botvinick. On the importance of single directions for generalization. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=rliuQjxCZ>.
- S. Na, Y. J. Choe, D.-H. Lee, and G. Kim. Discovery of natural language concepts in individual units of cnns. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=S1EERs09YQ>.
- A. Nguyen, A. Dosovitskiy, J. Yosinski, T. Brox, and J. Clune. Synthesizing the preferred inputs for neurons in neural networks via deep generator networks. In *Advances in Neural Information Processing Systems*, pages 3387–3395, 2016.


- A. Nguyen, J. Clune, Y. Bengio, A. Dosovitskiy, and J. Yosinski. Plug & play generative networks: Conditional iterative generation of images in latent space. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4467–4477, 2017.
- A. Nguyen, J. Yosinski, and J. Clune. Understanding neural networks via feature visualization: A survey. *arXiv preprint arXiv:1904.08939*, 2019.
- D. M. Powers. Recall & precision versus the bookmaker. In *Proceedings of the International Conference on Cognitive Science*. International Conference on Cognitive Science, 2003.
- A. Radford, R. Jozefowicz, and I. Sutskever. Learning to generate reviews and discovering sentiment. *arXiv preprint arXiv:1704.01444*, 2017.
- K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- K. Simonyan, A. Vedaldi, and A. Zisserman. Deep inside convolutional networks: Visualising image classification models and saliency maps. *arXiv preprint arXiv:1312.6034*, 2013.
- C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1–9, 2015.
- J. Yosinski, J. Clune, A. Nguyen, T. Fuchs, and H. Lipson. Understanding neural networks through deep visualization. *arXiv preprint arXiv:1506.06579*, 2015.
- M. D. Zeiler and R. Fergus. Visualizing and understanding convolutional networks. In *European conference on computer vision*, pages 818–833. Springer, 2014.
- B. Zhou, A. Khosla, A. Lapedriza, A. Oliva, and A. Torralba. Object detectors emerge in deep scene CNNs. In *International Conference on Learning Representations*, 2015.
- B. Zhou, D. Bau, A. Oliva, and A. Torralba. Interpreting deep visual representations via network dissection. *IEEE transactions on pattern analysis and machine intelligence*, 2018a.
- B. Zhou, Y. Sun, D. Bau, and A. Torralba. Revisiting the importance of individual units in CNNs via ablation. *arXiv preprint arXiv:1806.02891*, 2018b.

APPENDIX

5 METHODOLOGICAL DETAILS FOR THE BEHAVIORAL EXPERIMENT

One hundred generated images were made for every unit in layers conv5, fc6 and fc8 in AlexNet, as in Nguyen et al. (2017), and displayed as 10x10 image panels (figures 9 and Figures 7 and 8). A total of 3,299 image panels were used in the experiment (995 fc8, 256 conv5, and 2048 randomly selected fc6 image panels) and were divided into 64 counterbalanced lists of 51 or 52 (4 conv5, 15 or 16 fc8 and 32 fc6). 51 of the lists were assigned to 5 participants and 13 lists were assigned to 6 participants.

To test the interpretability of these units, paid volunteers were asked to look at image panels and asked if the images had an object / animal or place in common. If the answer was ‘yes’, they were asked to name that object simply (i.e. fish rather than goldfish). Analyses of common responses was carried out for any units where over 80% of humans agreed there was an object present, by reading the human responses and comparing them to both each other and to the output classes. Agreement was taken if the object was the same rough class. For example, ‘beer’, ‘glass’, and ‘drink’ were all considered to be in agreement in the general object of ‘drink’, and in agreement with both the classes of ‘wine glass’ and ‘beer’ as these classes were also general drink classes (this is an actual example, most responses were more obvious and required far less interpretation than that). Participants were given six practice trials, each with panels of 20 images before starting the main experiment. Practice trials included images that varied in their interpretability.



Q1. Can you identify multiple examples of an everyday object, place or animal?

Yes No

Click ‘Yes’ or ‘No’. In this example, the answer is ‘Yes’. After pressing ‘Yes’, two follow-up questions will appear, as shown below.

Q2. You answered ‘Yes’. Indicate what percentage of people you think would also answer ‘Yes’:

0-20% 20-40% 40-60% 60-80% 80-100%

Click on one of the boxes (in this case, 80-100% is an appropriate answer)

Q3. Please write which everyday object, place or animal you were referring to:

fish, goldfish

Type your answer in the text box. You may write a maximum of two things.

Figure 6: Example screen from the identification task shown to participants as part of the instructions. The images included on this practice trial are ImageNet2012 images, not AM images.

5.1 HUMAN INTERPRETATION OF ACTIVATION MAXIMIZATION IMAGES

Some examples of the 10x10 grids of activation maximisation images that were presented to participants are shown in Figures 7, 8 and 9. Figure 7 shows an example from conv5 that human participants agreed had no obvious object in common (although there are repeated shape motifs, the participants were specifically asked for objects, and not abstract concepts like shape or color. Figure 8

is also from the conv5 and was judged by participants as some images containing ‘dogs’. Figure 9 is the AM images for the supposed ‘butterfly detector’ unit example discussed in the paper.



Figure 7: Example activation maximisation images for unit conv5.65. These images were judged by humans to not contain any interpretable objects in common (although the reader may agree that there are some shape and colour similarities in the images).




Figure 8: Example activation maximisation images for unit conv5.183. These images were judged by humans to contain some interpretable images, in this case, of the type ‘dogs’.



Figure 9: Example activation maximisation images for unit fc6.1199. Whilst there are some butterfly wing shapes in these images, there are not obvious butterflies. N.B. the second highest activating class for this unit is ladybirds, and there are some orange round shapes that could conceivably be ladybug-alikes.

6 FURTHER DATA ON THE SELECTIVITY MEASURES

Further data on the selectivity measures across AlexNet. Table 2 demonstrates some highest values of CCMAS and precision in AlexNet.

Table 2: The units with the highest CCMAS and precision scores in AlexNet. Unit fc6.1199 was displayed in Fig. 3.

LAYER.UNIT	CCMAS	Precision
Top CCMAS units		
prob.322	0.991	100%
fc7.31	0.94	11%
fc6.582	0.93	1%
conv5.78	0.75	5%
Top precision units		
prob.0	0.99	100%
fc7.255	0.90	97%
fc6.1199	0.92	95%
conv5.0	0.55	77%

6.1 CORRELATION BETWEEN SELECTED MEASURES

	CCMAS	recall _{0.95}	Max. Inf.	No. classes in top100
precision	0.38	0.30	0.15	-0.68
CCMAS		0.09	0.14	-0.47
recall _{0.95}			0.10	-0.19
Max. Inf.				-0.22

Table 3: The correlations between the different measures. (All p 's < .001)

7 FURTHER ISSUES WITH THE CCMAS MEASURE

7.1 HISTORGRAMS AND DISTRIBUTION FITS FOR ACTIVATIONS IN UNIT fc6.1199

The CCMAS measure is based on comparing the mean activation of categories, and this is problematic for a few reasons. First, the majority of images do not activate a unit at all. For instance, our butterfly ‘detector’ unit fc6.1199 has 82.8% of images with an activation of 0.0 (see figure 10). This means that the CCMAS selectivity is largely determined by the proportion of images that have 0 activation rather than the distribution of images that strongly activate a unit. This leads to very different estimates of precision, top-class and localist selectivity that are concerned with the distribution of highly activated units.

The problem with the CCMAS measure when applied to the non-zero activations is that they are not normally distributed. As figure 11 demonstrates (for our example unit fc6.1199), the all non-zero activations follow an exponential curve and thus the mean is not a useful measure. Fitting a normal distribution to this data gives the blue normal distribution curve. Although the butterfly class can be roughly fit by a normal distribution, as the entire activations follow an exponential, the non-butterfly classes will be best fit by an exponential not a normal distribution. As the CCMAS requires a comparison of means, and the not-A classes follow an exponential, rather than normal distribution, it follows that the CCMAS will give misleading results.

Computing CCMAS on the basis of mean activations can produce highly non-intuitive as illustrated in figure 12 that plots three distributions of generated data from 10 classes of 100 points. We can see that the CCMAS gives the same (and maximal) score for the case where the unit is off to everything but a single item (figure 12a) as it does for a ‘grandmother cell’ unit (see figure 12c). The CCMAS




Figure 10: The distribution of activations on unit fc6.1199 for all images (left). 82.8% of activations are zero across all classes, only 5.8% of the butterfly class are zero. Unit fc6.1199 is a butterfly detector under Zhou et al. (2015)’s precision measure. Bins are 1.0 wide.

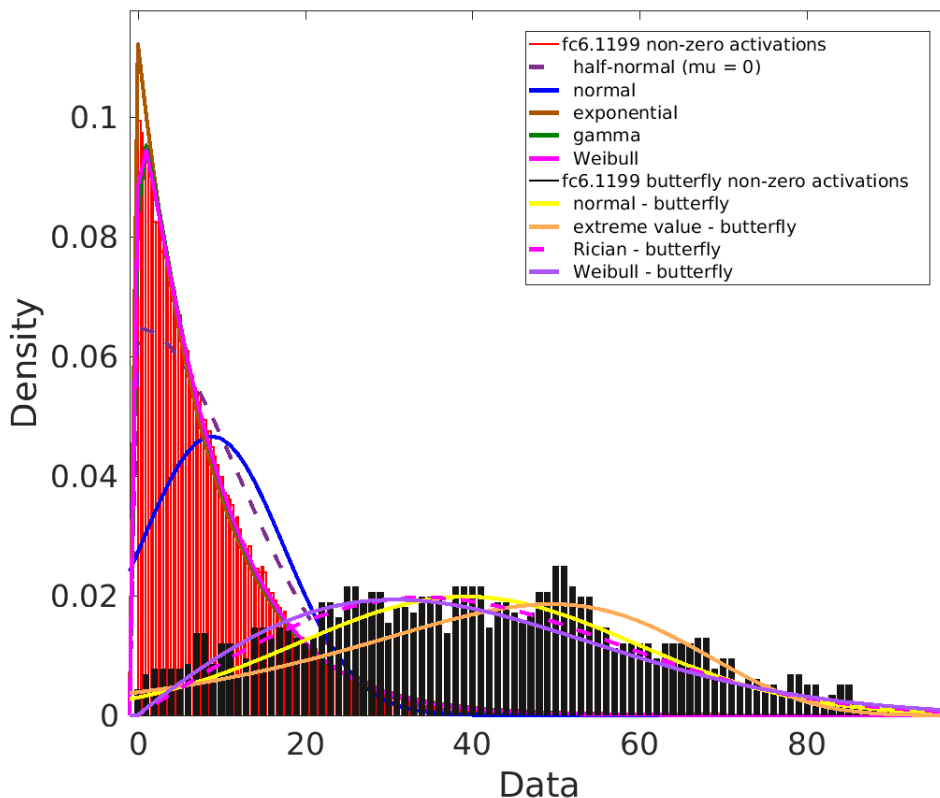


Figure 11: Fitting the non-zero activations for all classes (red) and the maximum activation class (black). For the superset of all the classes, the distribution is well-described by exponential-derived fits, normal-derived fits are bad. For the maximum activating class (butterfly), the distribution has a mean and can be well-described as a normal distribution.

also gives an incredibly low score when the means of A and not-A are similar, even if they perfectly separate the disjoint sets (figure 12b).

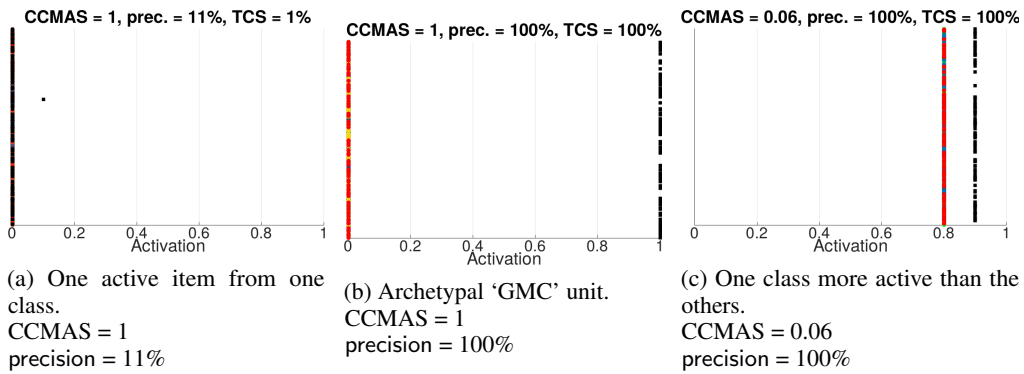


Figure 12: Example of where the CCMAS does not match intuitive understandings of selectivity. Generated example data: (a) If a unit is off to all but a single image from a large class of objects, the CCMAS for that class is 1 (maximum possible selectivity). (b) If a unit is strongly activated to all members of one class and off to everything else (an archetypal ‘grandmother’ cell) the CCMAS is the same as for (a) although the precision is vastly different. (c): If a unit has high activations for all classes, but one class (black squares) is 0.1 more than all others (coloured circles), the CCMAS is very low (0.06) despite being %100 precision. The generated examples are for 10 classes of 100 items

8 TESTING UNITS IN OTHER MODELS

To investigate units claimed by Zhou et al. (2018a) to be object detectors, we took units from the top layer and chose to use two categories that tended to have both a high number of object detectors and a high IoU for the concept: buses and cars. We used the first 100 images per class from the ImageNet 2012 dataset as our test data. Buses had three classes: ‘n04146614 school bus’, ‘n04487081 trolleybus, trolley coach, trackless trolley’, ‘n03769881 minibus’, and this corresponded to 300 items out of 100000 images. Cars: ‘n02930766 cab, hack, taxi, taxicab’, ‘n04037443 racer, race car, racing car’, ‘n03594945 jeep, landrover’, ‘n03770679 minivan’, ‘n02814533 beach wagon, station wagon, wagon, estate car, beach waggon, station waggon, waggon’, ‘n03670208 limousine, limo’ ‘n0377568 Model T’, ‘n03100240 convertible’, ‘n03459775 grille, radiator grille’, ‘n04285008 sports car, sport car’, ‘n04461696 tow truck, tow car, wrecker’, ‘n02965783 car mirror’, ‘n02974003 car wheel’, which is 1300 items. Data for all car and bus unit detectors for VGG trained on places 365 are shown in table 4.

unit	cat.	IoU	no. $a_x > 0$ $x \in A$	no. $a_x > 0$ $x \in \neg A$	μ_A	$\mu_{\neg A}$	precision	CCMAS
conv5_3 ₁₉₁	buses	0.15	99.0%	63.9%	131.9	16.1	0.45	0.783
conv5_3 ₂₀	buses	0.15	99.0%	49.1%	157.6	15.2	0.53	0.824
conv5_3 ₃₃₃	buses	0.08	99.0%	71.4%	101.7	17.5	0.24	0.705
conv5_3 ₁₄₅	buses	0.07	97.3%	61.7%	75.5	12.5	0.19	0.716
conv5_3 ₁₁₃	buses	0.06	97.4%	41.0%	62.8	9.1	0.12	0.747
conv5_3 ₄₄₃	buses	0.04	95.3%	38.2%	59.3	8.1	0.12	0.760
conv5_3 ₁₃₁	buses	0.04	93.7%	22.3%	54.0	5.86	0.08	0.804

Table 4: Selectivity measures for VGG-16, trained on Places-365, top convolutional layer units identified by Zhou et al. (2018a) as object detectors. Standard errors not shown for space, but were below ± 5 .

Representative statistics: Class A is 300 items, class not A is 99700 items. mean (A): 4.94 +/-0.38, std (A): 6.66. mean (NOT A): 4.01 +/-0.02, std (NOT A): 6.91. CCMAS(A):0.10

unit	cat.	IoU	correct	no. $a_x > 0$ $x \in A$	no. $a_x > 0$ $x \in \neg A$	μ_A	$\mu_{\neg A}$	precision	CCMAS
inception4e ₈₂₄	buses	0.17	T	100.0	91.4	41.0	11.8	0.27	0.553
inception4e ₇₄₅	buses	0.13	T	98.3	74.8	34.8	11.4	0.06	0.507
inception4e ₇₉₁	buses	0.11	T	98.3	71.4	32.7	5.3	0.41	0.719
inception4e ₁₉₄	buses	0.11	F	100.0	85.3	26.6	8.8	0.02	0.505
inception4e ₈₂	buses	0.11	T	100.0	97.3	26.7	10.9	0.14	0.420
inception4e ₇₃₆	buses	0.11	T	100.0	78.8	38.7	9.9	0.05	0.591
inception4e ₆₆₃	buses	0.10	F	96.0	38.0	33.4	3.7	0.15	0.802
inception4e ₉₄	buses	0.10	T	100.0	91.6	38.3	9.5	0.35	0.603
inception4e ₇₇₂	buses	0.08	F	97.3	54.6	21.7	5.2	0.0	0.610
inception4e ₁₁₃	buses	0.08	F	100.0	88.0	24.9	9.2	0.02	0.462
inception4e ₇₀₈	buses	0.06	F	100.0	85.1	29.7	9.1	0.02	0.529
inception4e ₈₀₁	buses	0.06	F	100.0	64.5	35.2	6.4	0.14	0.691
inception4e ₁₉₉	buses	0.06	F	99.7	92.2	21.5	7.7	0.09	0.473
inception4e ₈	buses	0.05	F	99.7	83.5	18.5	7.3	0.01	0.434
inception4e ₁₂₁	buses	0.05	F	100.0	90.4	17.9	8.9	0.01	0.336
inception4e ₆₂₂	buses	0.05	T	96.0	65.0	27.5	6.4	0.2	0.621
inception4e ₉₇	buses	0.04	T	99.3	86.4	21.1	9.3	0.04	0.389

Table 5: Selectivity measures for GoogLeNet, trained on Places-365, top convolutional layer units identified by Zhou et al. (2018a) as object detectors. Standard errors not shown for space, but were below ± 2 . A units is marked as correct if there was a single bus in the 4 example pictures on the website (http://netdissect.csail.mit.edu/dissect/googlenet_places365/), and false if not. Note that many object detectors pick out ‘bus like’ features in none bus objects.

unit	cat.	IoU	correct	no. $a_x > 0$ $x \in A$	no. $a_x > 0$ $x \in \neg A$	μ_A	$\mu_{\neg A}$	precision	CCMAS
inception4e494	buses	0.11	F	99.0	82.4	72.5	22.8	0.0	0.521
inception4e828	buses	0.10	T	100.0	72.6	109.4	17.6	0.45	0.723
inception4e569	buses	0.10	T	99.7	85.9	74.9	20.0	0.05	0.579
inception4e384	buses	0.10	T	100.0	71.6	67.0	18.5	0.0	0.566
inception4e455	buses	0.09	T	99.7	89.6	69.1	14.3	0.3	0.656
inception4e579	buses	0.09	T	100.0	97.0	91.5	26.0	0.23	0.557
inception4e331	buses	0.08	T	98.0	75.5	51.0	11.8	0.12	0.624
inception4e582	buses	0.08	T	100.0	83.4	125.7	21.95	0.58	0.703
inception4e498	buses	0.07	T	97.7	77.2	73.5	15.0	0.52	0.661
inception4e534	buses	0.07	F	99.3	81.2	62.7	19.	0.02	0.532
inception4e693	buses	0.07	T	98.7	91.2	75.4	22.3	0.15	0.544
inception4e673	buses	0.07	T	99.7	88.4	88.6	23.0	0.33	0.587
inception4e469	buses	0.06	T	98.7	78.1	34.7	14.6	0.0	0.409
inception4e207	buses	0.06	T	100.0	93.5	76.1	21.3	0.07	0.562
inception4e491	buses	0.06	F	99.0	74.5	41.1	13.7	0.01	0.501
inception4e645	buses	0.06	T	98.0	83.9	59.9	18.1	0.2	0.536
inception4e527	buses	0.06	F	100.0	91.5	58.0	21.7	0.0	0.456
inception4e511	buses	0.05	F	100.0	89.4	53.5	21.7	0.0	0.423
inception4e308	buses	0.05	F	100.0	89.4	53.5	21.7	0.0	0.423
inception4e541	buses	0.05	F	99.67	88.7	44.9	13.7	0.0	0.534
inception4e367	buses	0.05	T	97.3	80.3	37.7	15.4	0.02	0.420
inception4e665	buses	0.05	T	100.0	82.45	107.2	21.0	0.33	0.673
inception4e532	buses	0.05	T	100.0	91.5	52.9	22.4	0.05	0.405
inception4e297	buses	0.04	T	99.7	90.2	48.2	17.9	0.0	0.458
inception4e480	buses	0.04	T	100.0	92.9	69.4	21.4	0.02	0.529

Table 6: Selectivity measures for GoogLeNet, trained on ImageNet, top convolutional layer units identified by Zhou et al. (2018a) as object detectors. Standard errors not shown for space, but were below ± 2 . A units is marked as correct if there was a single bus in the 4 example pictures on the website (http://netdissect.csail.mit.edu/dissect/googlenet_imagenet/), and false if not. Note that many object detectors pick out 'bus like' features in none bus objects.

AUTHOR CONTRIBUTIONS

If you'd like to, you may include a section for author contributions as is done in many journals. This is optional and at the discretion of the authors.