

Interpretable CNNs for Object Classification

Quanshi Zhang¹, Xin Wang¹, Ying Nian Wu, Huilin Zhou, and Song-Chun Zhu²

Abstract—This paper proposes a generic method to learn interpretable convolutional filters in a deep convolutional neural network (CNN) for object classification, where each interpretable filter encodes features of a specific object part. Our method does not require additional annotations of object parts or textures for supervision. Instead, we use the same training data as traditional CNNs. Our method automatically assigns each interpretable filter in a high conv-layer with an object part of a certain category during the learning process. Such explicit knowledge representations in conv-layers of the CNN help people clarify the logic encoded in the CNN, i.e., answering what patterns the CNN extracts from an input image and uses for prediction. We have tested our method using different benchmark CNNs with various architectures to demonstrate the broad applicability of our method. Experiments have shown that our interpretable filters are much more semantically meaningful than traditional filters.

Index Terms—Convolutional neural networks, interpretable deep learning

1 INTRODUCTION

IN recent years, convolutional neural networks (CNNs) [9], [14], [17] have achieved superior performance in many visual tasks, such as object classification and detection. In spite of the good performance, a deep CNN has been considered a black-box model with weak feature interpretability for decades. Boosting the feature interpretability of a deep model gradually attracts increasing attention recently, but it presents significant challenges for state-of-the-art algorithms.

In this paper, we focus on a new task, i.e., without any additional annotations for supervision, revising a CNN to make its high conv-layers (e.g., the top two conv-layers) encode interpretable object-part knowledge. The revised CNN is termed an *interpretable CNN*.

More specifically, we propose a generic interpretable layer to construct the interpretable CNN. Each filter in the proposed interpretable layer automatically learns specific, discriminative object-part features. Filters in the interpretable layer are supposed to have some introspection of their feature representations and regularize their features towards object parts during the end-to-end learning. We trained interpretable CNNs on several benchmark datasets, and experimental results show that filters in the interpretable layer consistently represented the same object part across input images.

In addition, as discussed in [2], filters in low conv-layers usually describe textural patterns, while filters in high conv-layers are more likely to represent part patterns. Therefore, we focus on part-based interpretability and

propose a method to ensure each filter in a high conv-layer to represent an object part.

Note that our task of improving feature interpretability of a CNN is essentially different from the conventional visualization [6], [7], [20], [25], [29], [41] and diagnosis [2], [12], [19], [22] of pre-trained CNNs. The interpretable CNN learns more interpretable features, whereas previous methods mainly explain pre-trained neural networks, instead of learning new and more interpretable features.

Fig. 1 visualizes the difference between a traditional filter and our interpretable filter. In a traditional CNN, a filter usually describes a mixture of patterns. For example, the filter may be activated by both the head part and the leg part of a cat. In contrast, the filter in our interpretable CNN is expected to be activated by a certain part.

Thus, the goal of this study can be summarized as follows. We propose a generic interpretable conv-layer to construct the interpretable CNN. Feature representations of the interpretable conv-layer are interpretable, i.e., each filter in the interpretable conv-layer learns to consistently represent the same object part across different images. In addition, the interpretable conv-layer needs to satisfy the following properties:

- The interpretable CNN needs to be learned without any additional annotations of object parts for supervision. We use the same training samples as the original CNN for learning.
- The interpretable CNN does not change the loss function of the classification task, and it can be broadly applied to different benchmark CNNs with various architectures.
- As an exploratory research, learning strict representations of object parts may hurt a bit the discrimination power. However, we need to control the decrease within a small range.

Method: As shown in Fig. 2, we propose a simple yet effective loss. We simply add the loss to the feature map of each filter in a high conv-layer, so as to construct an interpretable conv-layer. The filter loss is proposed based on the

• Q. Zhang, X. Wang, and H. Zhou are with John Hopcroft Center, MoE Key Lab of Artificial Intelligence, AI Institute, Shanghai Jiao Tong University, Shanghai 200240, China. E-mail: {zqs1022, xin.wang, zhouhuilin116}@sjtu.edu.cn.

• Y.N. Wu and S.-C. Zhu are with the University of California, Los Angeles, Los Angeles, CA 90095 USA. E-mail: {ywu, sczhu}@stat.ucla.edu.

Manuscript received 8 Jan. 2019; revised 11 Jan. 2020; accepted 11 Mar. 2020. Date of publication 30 Mar. 2020; date of current version 2 Sept. 2021.

(Corresponding author: Quanshi Zhang.)

Recommended for acceptance by B. Ommer.

Digital Object Identifier no. 10.1109/TPAMI.2020.2982882

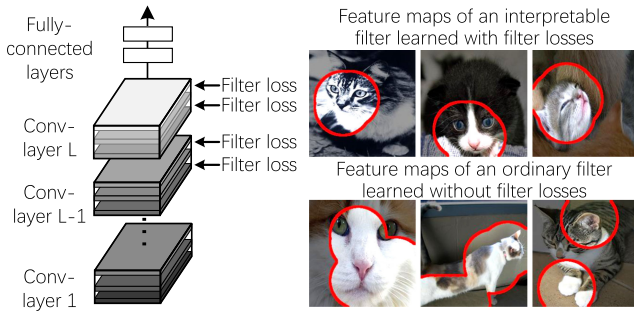


Fig. 1. Comparison of an interpretable filter's feature maps with a filter's feature maps in a traditional CNN.

assumption that only a single target object is contained in the input image. The filter loss forces the filter to encode the representation of a specific object part.

Theoretically, we can prove that the loss encourages a low entropy of inter-category activations and a low entropy of spatial distributions of neural activations. In other words, this loss ensures that (i) each filter must encode one object part of a single object category, instead of representing multiple categories; (ii) The feature must be consistently triggered by a single specific part across multiple images, rather than be simultaneously triggered by different object regions in each input image. It is assumed that repetitive patterns on different object regions are more likely to describe low-level textures, instead of high-level parts.

Value of Feature Interpretability: Such explicit object-part representations in conv-layers of CNN can help people clarify the decision-making logic encoded in the CNN at the object-part level. Given an input image, the interpretable conv-layer enables people to explicitly identify which object parts are memorized and used by the CNN for classification without ambiguity. Note that the automatically learned object part may not have an explicit name, e.g., a filter in the interpretable conv-layer may describe a partial region of a semantic part or the joint of two parts.

In critical applications, clear disentanglement of visual concepts in high conv-layers helps people trust a network's prediction. As analyzed in [46], a good performance on testing images cannot always ensure correct feature representations considering potential dataset bias. For example, in [46], a CNN used an unreliable context—eye features—to identify the “lipstick” attribute of a face image. Therefore, people need to semantically and visually explain what patterns are learned by the CNN.

Contributions: In this paper, we focus on a new task, i.e., learning an interpretable CNN without any part annotations in an end-to-end manner, where filters of high conv-layers represent specific object parts. We propose a simple yet effective method to learn interpretable filters, and the method can be broadly applied to different benchmark CNNs. Experiments show that our approach has significantly boosted feature interpretability of CNNs.

A preliminary version of this paper appeared in [47].

2 RELATED WORK

The interpretability and the discrimination power are two crucial aspects of a CNN [2]. In recent years, different methods are developed to explore the semantics hidden inside a




Fig. 2. Architectures of an ordinary conv-layer and an interpretable conv-layer. Solid and dashed lines indicate the forward and backward propagations, respectively. During the forward propagation, our CNN assigns each interpretable filter with a specific mask *w.r.t.* each input image during the learning process.

CNN. Our previous paper [49] provides a comprehensive survey of recent studies in exploring visual interpretability of neural networks, including (i) the visualization and diagnosis of CNN representations, (ii) approaches for disentangling CNN representations into graphs or trees, (iii) the learning of CNNs with disentangled and interpretable representations, and (iv) middle-to-end learning based on model interpretability.

2.1 Interpretation of Pre-Trained Neural Networks

Network Visualization: Visualization of filters in a CNN is the most direct way of exploring the pattern that is encoded by the filter. Gradient-based visualization [20], [29], [41] showed the appearance that maximized the score of a given unit. Furthermore, Bau *et al.* [2] defined and analyzed the interpretability of each filter. They classified all potential semantics into the following six types, *objects*, *parts*, *scenes*, *textures*, *materials*, and *colors*. We can further summarize the semantics of *objects* and *parts* as part patterns with specific contours and consider the other four semantics as textural patterns without explicit shapes. Recently, [21] provided tools to visualize filters of a CNN. Dosovitskiy *et al.* [6] proposed up-convolutional nets to invert feature maps of conv-layers to images. However, up-convolutional nets cannot mathematically ensure the visualization result reflects actual neural representations.

Although above studies can produce clear visualization results, theoretically, gradient-based visualization of a filter usually selectively visualizes the strongest activations of a filter in a high conv-layer, instead of illustrating knowledge hidden behind all activations of the filter; otherwise, the visualization result will be chaotic. Similarly, [2] selectively analyzed the semantics of the highest 0.5 percent activations of each filter. In comparisons, we aim to purify the semantic meaning of each filter in a high conv-layer, i.e., letting most activations of a filter be explainable, instead of extracting meaningful neural activations for visualization.

Pattern Retrieval: Unlike passive visualization, some methods actively retrieve certain units with certain meanings from CNNs. Just like mid-level features [31] of images, pattern retrieval mainly focuses on mid-level representations in

conv-layers. For example, Zhou *et al.* [50], [51] selected units from feature maps to describe “scenes”. Simon *et al.* discovered objects from feature maps of conv-layers [27], and selected certain filters to represent object parts [28]. Zhang *et al.* [43] extracted certain neural activations of a filter to represent object parts in a weakly-supervised manner. They also disentangled CNN representations via active question-answering and summarized the disentangled knowledge using an And-Or graph [44]. [45] used human interactions to refine the AOG representation of CNN knowledge. [8] used a gradient-based method to explain visual question-answering. Other studies [13], [18], [35], [37] selected filters or neural activations with specific meanings from CNNs for various applications. Unlike the retrieval of meaningful neural activations from noisy features, our method aims to substantially boost the interpretability of features in intermediate conv-layers.

Model Diagnosis: Many approaches have been proposed to diagnose CNN features, including exploring semantic meanings of convolutional filters [33], the examination of the transferability of filters [40], and the analysis of feature distributions of different categories [1]. The LIME [22] and the SHAP [19] are general methods to extract input units of a neural network that are used for the prediction score. For CNNs oriented to visual tasks, gradient-based visualization methods [7], [25] and [15] extracted image regions that are responsible for the network output, in order to clarify the logic of network prediction. These methods require people to manually check image regions accountable for the label prediction for each testing image. [11] extracted relationships between representations of various categories from a CNN. Zhang *et al.* [5] explained the success of knowledge distillation by quantifying knowledge points hidden in a CNN. In contrast, given an interpretable CNN, people can directly identify object parts or filters that are used for prediction.

As discussed by Zhang *et al.* [46], knowledge representations of a CNN may be significantly biased due to dataset bias, even though the CNN sometimes exhibits good performance. For example, a CNN may extract unreliable contextual features for prediction. Network-attack methods [12], [32], [33] diagnosed network representation flaws using adversarial samples of a CNN. For example, influence functions [12] can be used to generate adversarial samples, in order to fix the training set and further debug representations of a CNN. [16] discovered blind spots of knowledge representation of a pre-trained CNN in a weakly-supervised manner.

Distilling Neural Networks into Explainable Models: Furthermore, some methods distilled CNN knowledge into another model with interpretable features for explanations. [34] distilled knowledge of a neural network into an additive model to explain the knowledge inside the network. [48] roughly represented the rationale of each CNN prediction using a semantic tree structure. Each node in the tree represented a decision-making mode of the CNN. Similarly, [42] used a semantic graph to summarize and explain all part knowledge hidden inside conv-layers of a CNN.

2.2 Learning Interpretable Feature Representations

Unlike the diagnosis and visualization of pre-trained CNNs, some approaches were developed to learn meaningful feature representations in recent years. Automatically learning

interpretable feature representations without additional human annotations proposes new challenges to state-of-the-art algorithms. For example, [23] required people to label dimensions of the input that were related to each output, in order to learn a better model. Hu *et al.* [10] designed logic rules to regularize network outputs during the learning process. Sabour *et al.* [24] proposed a capsule model, where each feature dimension of a capsule may represent a specific meaning. Similarly, we invent a generic filter loss to regularize the representation of a filter to improve its interpretability.

In addition, unlike the visualization methods (e.g., the Grad-CAM method [25]) using a single saliency map for visualization, our interpretable CNN disentangles feature representations and uses different filters to represent the different object parts.

3 ALGORITHM

Given a target conv-layer of a CNN, we expect each filter in the conv-layer to be activated by a certain object part of a certain category, while remain inactivated on images of other categories.¹ Let \mathbf{I} denote a set of training images, where $\mathbf{I}_c \subset \mathbf{I}$ represents the subset that belongs to category c , ($c = 1, 2, \dots, C$). Theoretically, we can use different types of losses to learn CNNs for multi-class classification and binary classification of a single class (i.e., $c = 1$ for images of a category and $c = 2$ for random images).

In the following paragraphs, we focus on the learning of a single filter f in a conv-layer. Fig. 2 shows the architecture of our interpretable conv-layer. We add a loss to the feature map x of the filter f after the ReLU operation. The filter loss $Loss_f$ pushes the filter f to represent a specific object part of the category c and keep silent on images of other categories. Please see Section 3.2 for the determination of the category c for the filter f . Let $\mathbf{X} = \{x | x = f(I) \in \mathbb{R}^{n \times n}, I \in \mathbf{I}\}$ denote a set of feature maps of f after a ReLU operation *w.r.t.* different images. Given an input image $I \in \mathbf{I}_c$, the feature map in an intermediate layer $x = f(I)$ is an $n \times n$ matrix, $x_{ij} \geq 0$. If the target part appears, we expect the feature map $x = f(I)$ to exclusively activate at the target part’s location; otherwise, the feature map should keep inactivated.

Therefore, a high interpretability of the filter f requires a high mutual information between the feature map $x = f(I)$ and the part location, i.e., the part location can roughly determine activations on the feature map x .

Accordingly, we formulate the filter loss as the minus mutual information, as follows.

$$Loss_f = -MI(\mathbf{X}; \mathbf{\Omega}) = - \sum_{\mu \in \mathbf{\Omega}} p(\mu) \sum_x p(x|\mu) \log \frac{p(x|\mu)}{p(x)}, \quad (1)$$

where $MI(\cdot)$ denotes the mutual information; $\mathbf{\Omega} = \{\mu_1, \mu_2, \dots, \mu_{n^2}\} \cup \{\mu^-\}$. We use $\mu_1, \mu_2, \dots, \mu_{n^2}$ to denote the n^2 neural units on the feature map x , each $\mu = [i, j] \in \mathbf{\Omega}$, $1 \leq i, j \leq n$, corresponding to a location candidate for the target part. μ^- denotes a dummy location for the case when the target part does not appear on the image.

1. To avoid ambiguity, we evaluate or visualize the semantic meaning of each filter by using the feature map after the ReLU and mask operations.




Fig. 3. Templates of T_{μ_i} . We show a toy example of $n = 3$. Each template T_{μ_i} matches to a feature map x when the target part mainly triggers the i th unit in x . In fact, the algorithm also supports a round template based on the L-2 norm distance. Here, we use the L-1 norm distance instead to speed up the computation.

Given an input image, the above loss forces each filter to match and only match one of the templates, i.e., making the feature map of the filter contain a single significant activation peak at most. This ensures each filter to represent a specific object part.

- $p(\mu)$ measures the probability of the target part appearing at the location μ . If annotations of part locations are given, then the computation of $p(\mu)$ is simple. People can manually assign a semantic part with the filter f , and then $p(\mu)$ can be determined using part annotations.

However, in our study, the target part of filter f is not pre-defined before the learning process. Instead, the part corresponding to f needs to be determined during the learning process. More crucially, we do not have any ground-truth annotations of the target part, which boosts the difficulty of calculating $p(\mu)$.

- The conditional likelihood $p(x|\mu)$ measures the fitness between a feature map x and the part location $\mu \in \Omega$. In order to simplify the computation of $p(x|\mu)$, we design n^2 templates for f , $\{T_{\mu_1}, T_{\mu_2}, \dots, T_{\mu_{n^2}}\}$. As shown in Fig. 3, each template T_{μ_i} is an $n \times n$ matrix. T_{μ_i} describes the ideal distribution of activations for the feature map x when the target part mainly triggers the i th unit in x . In addition, we also design a negative template T^- corresponding to the dummy location μ^- . The feature map can match to T^- , when the target part does not appear on the input image. In this study, the prior probability is given as $p(\mu_i) = \frac{\alpha}{n^2}$, $p(\mu^-) = 1 - \alpha$, where α is a constant prior likelihood.

Note that in Equation (1), we do not manually assign filters with different categories. Instead, we use the negative template μ^- to help the assignment of filters. I.e., the negative template ensures that each filter represents a specific object part (if the input image does not belong to the target part, then the input image is supposed to match μ^-), which also ensures a clear assignment of filters to categories. Here, we assume two categories do not share object parts, e.g., eyes of dogs and those of cats do not have similar contextual appearance.

We define $p(x|\mu)$ below, which follows a standard form widely used in [26], [39].

$$p(x|\mu) \approx p(x|T_{\mu}) = \frac{1}{Z_{\mu}} \exp[\text{tr}(x \cdot T_{\mu})], \quad (2)$$

where $Z_{\mu} = \sum_{x \in \mathcal{X}} \exp[\text{tr}(x \cdot T_{\mu})]$. $\text{tr}(\cdot)$ indicates the trace of a matrix, and $\text{tr}(x \cdot T_{\mu}) = \sum_{ij} x_{ij} t_{ji}$, $x, T_{\mu} \in \mathbb{R}^{n \times n}$. $p(x) = \sum_{\mu} p(\mu)p(x|\mu)$.

Part Templates. As shown in Fig. 3, a negative template is given as $T^- = (t_{ij}^-)$, $t_{ij}^- = -\tau < 0$, where τ is a positive constant. A positive template corresponding to μ is given as $T_{\mu} = (t_{ij}^+)$, $t_{ij}^+ = \tau \cdot \max(1 - \beta \frac{\| [i,j] - \mu \|_1}{n}, -1)$, where $\|\cdot\|_1$ denotes the L-1 norm distance. Note that the lowest value in a positive template is -1 instead of 0. It is because that the negative value in the template penalizes neural activations outside the domain of the highest activation peak, which ensures each filter mainly has at most a single significant activation peak.

3.1 Part Localization and the Mask Layer

Given an input image I , the filter f computes a feature map x after the ReLU operation. Without ground-truth annotations of the target part for f , in this study, we determine the part location on x during the learning process. We consider the neural unit with the strongest activation $\hat{\mu} = \arg\max_{\mu=[i,j]} x_{ij}$, $1 \leq i, j \leq n$ as the target part location.

As shown in Fig. 2, we add a mask layer above the interpretable conv-layer. Based on the estimated part position $\hat{\mu}$, the mask layer assigns a specific mask with x to filter out noisy activations. The mask operation is separate from the filter loss in Equation (1). Our method selects the template $T_{\hat{\mu}}$ w.r.t. the part location $\hat{\mu}$ as the mask. We compute $x^{\text{masked}} = \max\{x \circ T_{\hat{\mu}}, 0\}$ as the output masked feature map, where \circ denotes the Hadamard (element-wise) product. The mask operation supports gradient back-propagations.

Fig. 4 visualizes the masks $T_{\hat{\mu}}$ chosen for different images, and compares the original and masked feature maps. The CNN selects different templates for different images.

Note that although a filter usually has much stronger neural activations on the target category than on other categories, the magnitude of neural activations is still not discriminative enough for classification. Moreover, during the testing process, people do not have ground-truth class labels of input images. Thus, to ensure stable feature extraction, our method only selects masks from the n^2 positive templates $\{T_{\mu_i}\}$ and omits the negative template T^- for all images, no matter whether or not input images contain the target part. Such an operation is conducted during the forward process for both training and testing processes.

3.2 Learning

We train the interpretable CNN in an end-to-end manner. During the forward-propagation process, each filter in the CNN passes its information in a bottom-up manner, just like traditional CNNs. During the back-propagation, each filter in an interpretable conv-layer receives gradients w.r.t. its feature map x from both the final task loss $\mathcal{L}(\hat{y}_k, y_k^*)$ on the k th sample and the filter loss, Loss_f , as follows:

$$\frac{\partial \text{Loss}}{\partial x_{ij}} = \lambda \sum_f \frac{\partial \text{Loss}_f}{\partial x_{ij}} + \frac{1}{N} \sum_{k=1}^N \frac{\partial \mathcal{L}(\hat{y}_k, y_k^*)}{\partial x_{ij}}, \quad (3)$$

where λ is a weight. Then, we back propagate $\frac{\partial \text{Loss}}{\partial x_{ij}}$ to lower layers and compute gradients w.r.t. feature maps and gradients w.r.t. parameters in lower layers to update the CNN.

For implementation, gradients of Loss_f w.r.t. each element x_{ij} of feature map x are computed as follows.

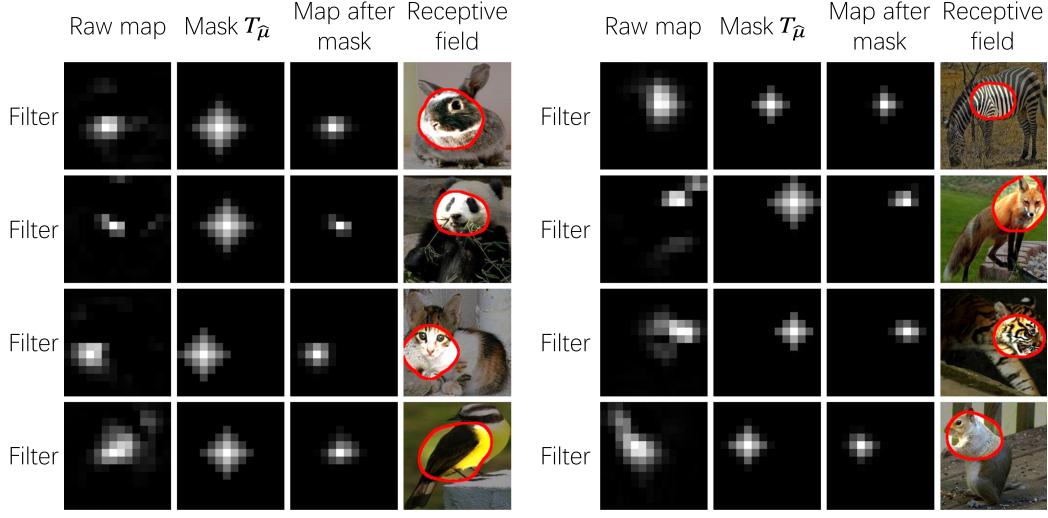


Fig. 4. Given an input image I , from left to right, we consequently show the feature map of a filter after the ReLU layer x , the assigned mask $T_{\hat{\mu}}$, the masked feature map x^{masked} , and the image-resolution RF of activations in x^{masked} computed by [50].

$$\begin{aligned} \frac{\partial \text{Loss}_f}{\partial x_{ij}} &= \frac{1}{Z_{\mu}} \sum_{\mu} p(\mu) t_{ij} e^{tr(x \cdot T_{\mu})} \left\{ tr(x \cdot T_{\mu}) - \log [Z_{\mu} p(x)] \right\} \\ &\approx \frac{p(\hat{\mu}) \hat{t}_{ij}}{Z_{\hat{\mu}}} e^{tr(x \cdot T_{\hat{\mu}})} \left\{ tr(x \cdot T_{\hat{\mu}}) - \log Z_{\hat{\mu}} - \log p(x) \right\}, \end{aligned} \quad (4)$$

where $T_{\hat{\mu}}$ is the target template for feature map x . If the input image I belongs to the target category of filter f , then $\hat{\mu} = \arg\max_{\mu=[i,j]} x_{ij}$. If image I belongs to other categories, then $\hat{\mu} = \mu^-$. Please see the appendix for the proof of the above equation, which can be found on the Computer Society Digital Library at <http://doi.ieeecomputersociety.org/10.1109/TPAMI.2020.2982882>.

Considering $\forall \mu \in \Omega \setminus \{\hat{\mu}\}, e^{tr(x \cdot T_{\hat{\mu}})} \gg e^{tr(x \cdot T_{\mu})}$ and $p(\hat{\mu}) \gg p(\mu)$ after initial learning episodes, we make the above approximation to simplify the computation. Because $Z_{\hat{\mu}}$ is computed using numerous feature maps, we can roughly treat $Z_{\hat{\mu}}$ as a constant to compute gradients in the above equation. We gradually update the value of $Z_{\hat{\mu}}$ during the training process. More specifically, we can use a subset of feature maps to approximate the value of Z_{μ} , and continue to update Z_{μ} when we receive more feature maps during the training process. Similarly, we can approximate $p(x) = \sum_{\mu} p(\mu) p(x|\mu) = \sum_{\mu} p(\mu) \frac{\exp[tr(x \cdot T_{\mu})]}{Z_{\mu}} \approx \sum_{\mu} p(\mu) \mathbb{E}_x \left[\frac{\exp[tr(x \cdot T_{\mu})]}{Z_{\mu}} \right]$ without huge computation.

Determining the Target Category for Each Filter. We need to assign each filter f with a target category \hat{c} to approximate gradients in Equation (4). We simply assign the filter f with the category \hat{c} whose images activate f the most, i.e., $\hat{c} = \arg\max_c \mathbb{E}_{x=f(I): I \in \mathcal{I}_c} \sum_{ij} x_{ij}$.

3.3 Understanding the Filter Loss

The filter loss in Equation (1) can be re-written as

$$\text{Loss}_f = -H(\Omega) + H(\Omega'|\mathbf{X}) + \sum_x p(\Omega^+, x) H(\Omega^+|X=x), \quad (5)$$

where $\Omega' = \{\mu^-, \Omega^+\}$. The first term $H(\Omega) = -\sum_{\mu \in \Omega} p(\mu) \log p(\mu)$ is a constant prior entropy of part locations.

Thus, the filter loss minimizes two conditional entropies, $H(\Omega'|\mathbf{X})$ and $H(\Omega^+|X=x)$. Please see the appendix for the proof of the above equation, available in the online supplemental material.

- *Low inter-category entropy:* The second term $H(\Omega' = \{\mu^-, \Omega^+\}|\mathbf{X})$ is computed as

$$H(\Omega' = \{\mu^-, \Omega^+\}|\mathbf{X}) = - \sum_x p(x) \sum_{\mu \in \{\mu^-, \Omega^+\}} p(\mu|x) \log p(\mu|x), \quad (6)$$

where $\Omega^+ = \{\mu_1, \dots, \mu_{n^2}\} \subset \Omega$, $p(\Omega^+|x) = \sum_{\mu \in \Omega^+} p(\mu|x)$. We define the set of all real locations Ω^+ as a single label to represent category c . We use the dummy location μ^- to roughly indicate matches to other categories.

This term encourages a low conditional entropy of inter-category activations, i.e., a well-learned filter f needs to be exclusively activated by a certain category c and keep silent on other categories. We can use a feature map x of f to identify whether or not the input image belongs to category c , i.e., x fitting to either $T_{\hat{\mu}}$ or T^- , without significant uncertainty.

- *Low spatial entropy:* The third term in Equation (5) is given as

$$H(\Omega^+|X=x) = \sum_{\mu \in \Omega^+} \tilde{p}(\mu|x) \log \tilde{p}(\mu|x), \quad (7)$$

where $\tilde{p}(\mu|x) = \frac{p(\mu|x)}{p(\Omega^+|x)}$. This term encourages a low conditional entropy of the spatial distribution of x 's activations. I.e., given an image $I \in \mathcal{I}_c$, a well-learned filter should only be activated in a single region $\hat{\mu}$ of the feature map x , instead of being repetitively triggered at different locations.

4 EXPERIMENTS

In experiments, we applied our method to modify four types of CNNs with various architectures into interpretable CNNs and learned interpretable CNNs based on three benchmark datasets, in order to demonstrate the broad applicability. We learned interpretable CNNs for binary classification of a single category and multi-category classification. We used different techniques to visualize the

knowledge encoded in interpretable filters, in order to qualitatively illustrate semantic meanings of these filters. Furthermore, we used two types of evaluation metrics, i.e., the object-part interpretability and the location instability, to measure the semantic clarity of a filter.

Our experiments showed that an interpretable filter in our interpretable CNN usually consistently represented the same part through different input images, while a filter in an ordinary CNN mainly described a mixture of semantics.

We chose three benchmark datasets with part annotations for training and testing, including the ILSVRC 2013 DET Animal-Part dataset [43], the CUB200-2011 dataset [36], and the VOC Part dataset [4]. These datasets provide ground-truth bounding boxes of entire objects. For landmark annotations, the ILSVRC 2013 DET Animal-Part dataset [43] contains ground-truth bounding boxes of heads and legs of 30 animal categories. The CUB200-2011 dataset [36] contains a total of 11.8K bird images of 200 species, and the dataset provides center positions of 15 bird landmarks. The VOC Part dataset [4] contains ground-truth part segmentations of 107 object landmarks in six animal categories.

These three datasets contain ground-truth annotations of object landmarks² (parts). We used ground-truth annotations to evaluate the semantic clarity of each filter. As mentioned in [4], [43], animals usually consist of non-rigid parts, which present considerable challenges for part localization. As in [4], [43], we selected animal categories in the three datasets for testing.

We learned interpretable filters based on architectures of four typical CNNs for evaluation, including the AlexNet [14], the VGG-M [30], the VGG-S [30], the VGG-16 [30]. Note that skip connections in residual networks [9] make a single feature map contain patterns of different filters. Thus, we did not use residual networks for testing to simplify the evaluation. Given a CNN, all filters in the top conv-layer were set as interpretable filters. Then, we inserted another conv-layer with M filters above the top conv-layer, which did not change the size of output feature maps. I.e., we set $M = 512$ for the VGG-16, VGG-M, and VGG-S networks, and $M = 256$ for the AlexNet. Filters in the new conv-layer were also interpretable filters. Each filter was a $3 \times 3 \times M$ tensor with a bias term.

Implementation Details. We set parameters as $\tau = \frac{0.5}{n^2}$, $\alpha = \frac{n^2}{1+n^2}$, and $\beta \approx 4$. β was updated during the learning process. We set a decreasing weight for filter losses, i.e., $\lambda \propto \frac{1}{t} \mathbb{E}_{x \in \mathcal{X}} \max_{i,j} x_{ij}$ for the t th epoch. We initialized fully-connected (FC) layers and the new conv-layer, but we loaded parameters of the lower conv-layers from a CNN that was pre-trained using [14], [30]. We then fine-tuned parameters of all layers in the interpretable CNN using training images in the dataset. To enable a fair comparison, when we learned the traditional CNN as a baseline, we also initialized FC layers of the traditional CNN, used pre-trained parameters in conv-layers, and then fine-tuned the CNN.

4.1 Experiments

Binary Classification of a Single Category: We learned interpretable CNNs based on above four types of network architectures to classify each animal category in above three

datasets. We also learned ordinary CNNs using the same data for comparison. We used the logistic log loss for binary classification of a single category from random images. We followed experimental settings in [42], [43] to crop objects of the target category as positive samples. Images of other categories were regarded as negative samples.

Multi-Category Classification: We learned interpretable CNNs to classify the six animal categories in the VOC Part dataset [4] and also learned interpretable CNNs to classify the thirty categories in the ILSVRC 2013 DET Animal-Part dataset [43]. In experiments, we tried both the softmax log loss and the logistic log loss³ for multi-category classification.

4.2 Qualitative Visualization of Filters

We followed the method proposed by Zhou *et al.* [50] to compute the receptive fields (RFs) of neural activations of a filter. We used neural activations after ReLU and mask operations and scaled up RFs to the image resolution. As discussed in [2], the traditional idea of directly propagating the theoretical receptive field of a neural unit in a feature map back to the image plane cannot accurately reflect the real image-resolution RF of the neural unit (i.e., the image region that contributes most to the score of the neural unit). Therefore, we used the method of [50] to compute real RFs.

Studies in both [50] and [2] have introduced methods to compute real RFs of neural activations on a given feature map. For ordinary CNNs, we simply used a round RF for each neural activation. We overlapped all activated RFs in a feature map to compute the final RF of the feature map.

Fig. 5 shows RFs⁴ of filters in top conv-layers of CNNs, which were trained for binary classification of a single category. Each filter in interpretable CNNs was mainly activated by a certain object part, whereas feature maps of ordinary CNNs after ReLU operations usually represented various object parts and textures. The clear disentanglement of object-part representations can help people to quantify the contribution of different object parts to the network prediction. Fig. 5 shows the explanation for part contribution, which was generated by the method of [3].

We found that interpretable CNNs usually encoded head patterns of animals in its top conv-layer for classification, although no part annotations were used to train the CNN. We can understand such results from the perspective of the information bottleneck [38] as follows. (i) Our interpretable filters selectively encode the most distinct parts of each category (i.e., the head for most categories), which minimizes the conditional entropy of the final classification given feature maps of a conv-layer. (ii) Each interpretable filter represents a specific part of an object, which minimizes the mutual information between the input image and middle-layer feature maps. The interpretable CNN “forgets” as much irrelevant information as possible.

In addition to the visualization of RFs, we also visualized heatmaps for part distributions and Grad-CAM attention maps of an interpretable conv-layer. Fig. 6 shows heatmaps for distributions of object parts that were encoded in interpretable filters. Fig. 7 compares Grad-CAM visualizations [25]

2. To avoid ambiguity, a landmark is referred to as the *central position* of a semantic part (a part with an explicit name, e.g., a head, a tail). In contrast, the part corresponding to a filter does not have an explicit name.

3. We considered the output y_c for each category c independent to outputs for other categories, thereby a CNN making multiple independent binary classifications of different categories for each image. Table 7 reported the average accuracy of the multiple classification outputs of an image.

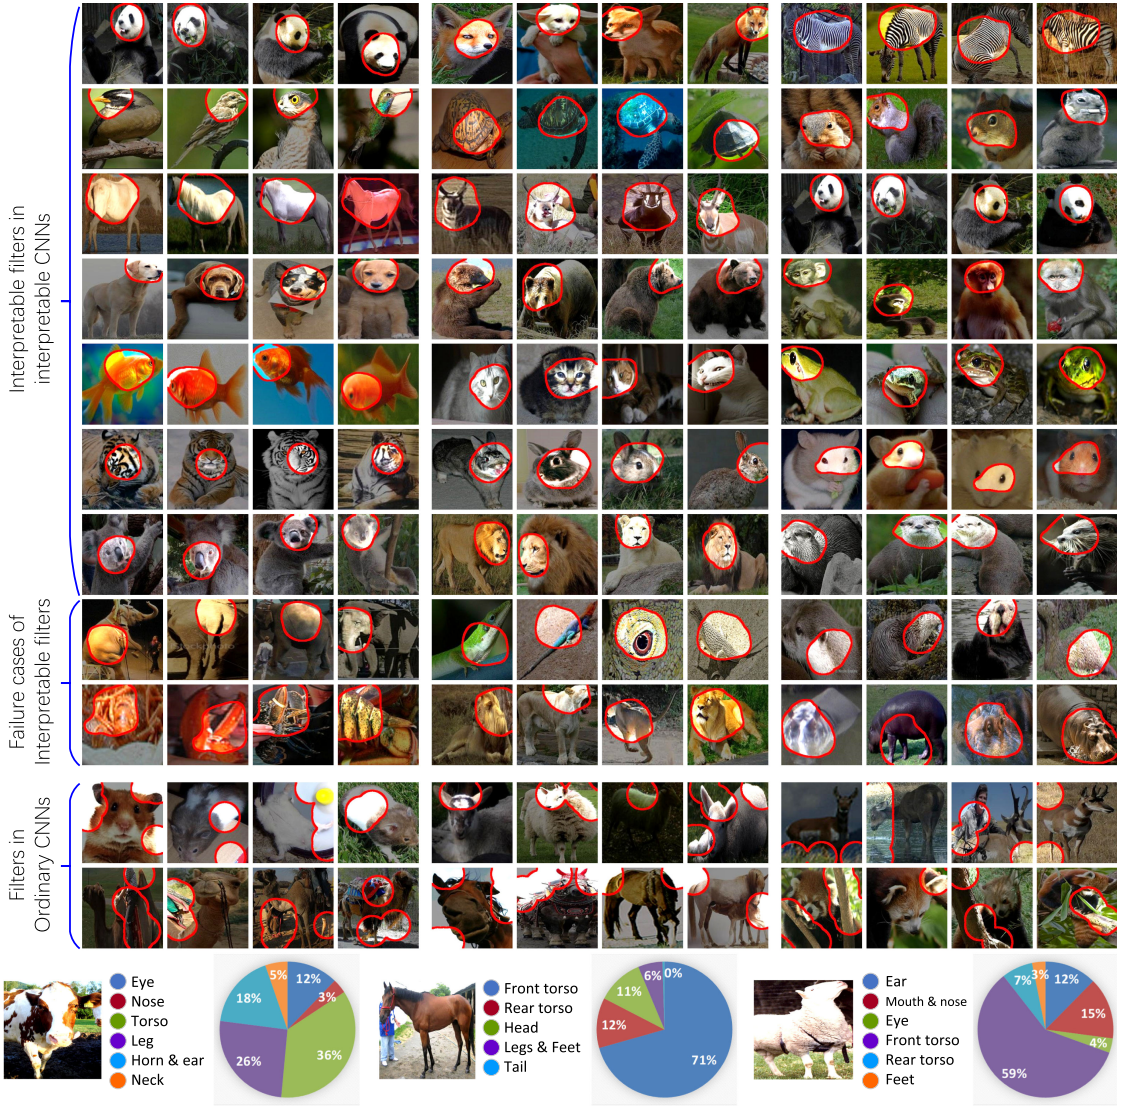


Fig. 5. Visualization of filters in top conv-layers (top) and quantitative contribution of object parts to the prediction (bottom). (top) We used [50] to estimate the image-resolution receptive field of activations in a feature map to visualize a filter's semantics. Each group of four feature maps for a category are computed using the same interpretable filter. These images show that each interpretable filter is consistently activated by the same object part through different images. Nine rows visualize filters in interpretable CNNs, and the last two rows correspond to filters in ordinary CNNs. (bottom) The clear disentanglement of object-part representations help people to quantify the contribution of different object parts to the network prediction. We show the explanation for part contribution, which was generated by the method of [3].

of an interpretable conv-layer and those of a traditional conv-layer. We chose the top conv-layer of the traditional VGG-16 net and the top conv-layer of the interpretable VGG-16 net for visualization. Interpretable filters usually selectively modeled distinct object parts of a category and ignored other parts.

4.3 Quantitative Evaluation of Part Interpretability

Filters in low conv-layers usually represent simple patterns or object details, whereas those in high conv-layers are more likely to describe large parts. Therefore, in experiments, we used the following two metrics to evaluate the semantic clarity of filters in the top conv-layer of a CNN.

4.3.1 Evaluation Metric: Part Interpretability

The metric was originally proposed by Bau *et al.* [2] to measure the object-part interpretability of filters. For each filter

f , \mathbf{X} denotes a set of feature maps after ReLU/mask operations on different input images. Then, the distribution of activation scores over all positions in all feature maps was computed. [2] set a threshold T_f such that $p(x_{ij} > T_f) = 0.005$ to select strongest activations from all positions $[i, j]$ from $x \in \mathbf{X}$ as valid activations for f 's semantics.

Then, image-resolution RFs of valid neural activations of each input image I were computed.⁴ The RFs on image I , termed S_f^I , corresponded to part regions of f .

4. [50] computes the RF when the filter represents an object part. Fig. 5 used RFs computed by [50] to visualize filters. However, when a filter in an ordinary CNN does not have consistent contours, it is difficult for [50] to align different images to compute an average RF. Thus, for ordinary CNNs, we simply used a round RF for each valid activation. We overlapped all activated RFs in a feature map to compute the final RF as mentioned in [2]. For a fair comparison, in Section 4.3.1, we uniformly applied these RFs to both interpretable CNNs and ordinary CNNs.

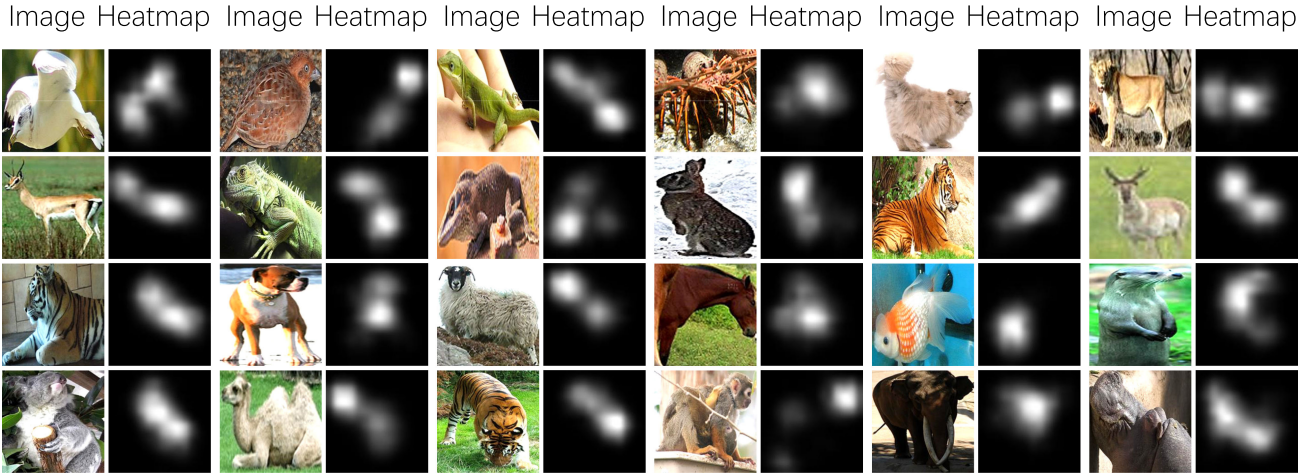


Fig. 6. Heatmaps for distributions of object parts that are encoded in interpretable filters. We use all filters in the top conv-layer to compute the heatmap. Interpretable filters usually selectively modeled distinct object parts of a category and ignored other parts.

The fitness between the filter f and the k th part on image I was reported as the intersection-over-union score $IoU_{f,k}^I = \frac{\|S_f^I \cap S_k^I\|}{\|S_f^I \cup S_k^I\|}$ where S_k^I represents the ground-truth mask of the k th part on image I . Given an image I , the filter f was associated with the k th part if $IoU_{f,k}^I > 0.2$. The criterion $IoU_{f,k}^I > 0.2$ was stricter than $IoU_{f,k}^I > 0.04$ in [2], because object-part semantics usually needs a stricter criterion than textural semantics and color semantics in [2]. The average probability of the k th part being associating with the filter f was reported as $P_{f,k} = \mathbb{E}_{I: \text{withk-thpart}} \mathbf{1}(IoU_{f,k}^I > 0.2)$. Note that a single filter may be associated with multiple object parts in an image. The highest probability of part association for each filter was used as the interpretability of filter f , i.e., $P_f = \max_k P_{f,k}$.

For the binary classification of a single category, we used testing images of the target category to evaluate the feature interpretability. In the VOC Part dataset [4], four parts were chosen for the *bird* category. We merged segments of the

head, beak, and l/r-eyes as the head part, merged segments of the torso, neck, and l/r-wings as the torso part, merged segments of l/r-legs/feet as the leg part, and used the tail segment as the fourth part. We used five parts for both the *cat* category and the *dog* category. We merged segments of the head, l/r-eyes, l/r-ears, and nose as the head part, merged segments of the torso and neck as the torso part, merged segments of frontal l/r-legs/paws as the frontal legs, merged segments of back l/r-legs/paws as the back legs, and used the tail as the fifth part. Part definitions for the *cow*, *horse*, and *sheep* category were similar those for the *cat* category, except for that we omitted the tail part of these categories. In particular, we added l/r-horn segments of the horse to the head part. The average part interpretability P_f over all filters was computed for evaluation.

For the multi-category classification, we first determined the target category \hat{c} for each filter f i.e., $\hat{c} = \arg\max_c \mathbb{E}_{x=f(I): I \in \mathcal{I}_c} \sum_{i,j} x_{ij}$. Then, we computed f 's object-part interpretability using images of the target category \hat{c} by following above instructions.

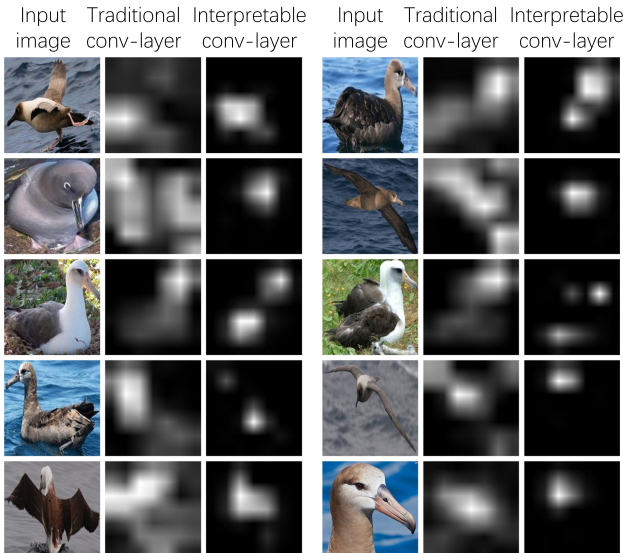


Fig. 7. Grad-CAM visualizations [25] of the traditional conv-layer and the interpretable conv-layer. Unlike the traditional conv-layer, the interpretable conv-layer usually selectively modeled distinct object parts of a category and ignored other parts.

4.3.2 Evaluation Metric: Location Instability

The second metric measures the instability of part locations, which was used in [42], [47]. It is assumed that if f consistently represented the same object part through different objects, then distances between the inferred part $\hat{\mu}$ and some ground-truth landmarks² should keep stable among different objects. For example, if f represented the shoulder part without ambiguity, then the distance between the inferred position and the head will not change a lot among different objects.

Therefore, the deviation of the distance between the inferred position $\hat{\mu}$ and a specific ground-truth landmark among different images was computed. The location $\hat{\mu}$ was inferred as the neural unit with the highest activation on f 's feature map. We reported the average deviation *w.r.t.* different landmarks as the location instability of f .

Please see Fig. 8. Given an input image I , $d_I(p_k, \hat{\mu}) = \frac{\|p_k - p(\hat{\mu})\|}{\sqrt{w^2 + h^2}}$ denotes the normalized distance between the inferred part and the k th landmark p_k , where $p(\hat{\mu})$ is




Fig. 8. Notation for computing the location instability.

referred to as the center of the unit $\hat{\mu}$'s RF. $\sqrt{w^2 + h^2}$ measures the diagonal length of I . $D_{f,k} = \sqrt{\text{var}_I[d_I(p_k, \hat{\mu})]}$ is termed as the *relative location deviation* of filter f w.r.t. the k th landmark, where $\text{var}_I[d_I(p_k, \hat{\mu})]$ is the variation of $d_I(p_k, \hat{\mu})$. Because each landmark could not appear in all testing images, for each filter f , the metric only used inference results on top-ranked 100 images with the highest inference scores to compute $D_{f,k}$. In this way, the average of relative location deviations of all the filters in a conv-layer w.r.t. all K landmarks, i.e., $\mathbb{E}_f \mathbb{E}_{k=1}^K D_{f,k}$, was reported as the location instability of f .

We used the most frequent object parts as landmarks to measure the location instability. For the ILSVRC 2013 DET Animal-Part dataset [43], we used the *head* and *frontal legs* of each category as landmarks for evaluation. For the VOC Part dataset [4], we selected the *head*, *neck*, and *torso* of each category as landmarks. For the CUB200-2011 dataset [36], we used the *head*, *back*, *tail* of birds as landmarks.

In particular, for multi-category classification, we first determined the target category of for each filter f and then computed the relative location deviation $D_{f,k}$ using landmarks of f 's target category. Because filters in baseline CNNs did not exclusively represent a single category, we simply assigned filter f with the category whose landmarks can achieve the lowest location deviation to simplify the computation. I.e., for a baseline CNN, we used $\mathbb{E}_f \min_c \mathbb{E}_{k \in \text{Part}_c} D_{f,k}$ to evaluate the location instability, where Part_c denotes the set of part indexes belonging to category c .

4.3.3 Comparisons Between Metrics of Filter Interpretability and Location Instability

Although the filter interpretability [2] and the location instability [47] are the two most state-of-the-art metrics to evaluate the interpretability of a convolution filter, these metrics still have some limitations.

First, the filter interpretability [2] assumes that the feature map of an automatically learned filter should well match the ground-truth segment of a semantic part (with an explicit part name), an object, or a texture. For example, it assumes that a filter may represent the exact segment of the head part. However, without ground-truth annotations of object parts or textures for supervision, there is no mechanism to assign explicit semantic meanings with filters during the learning process. In most cases, filters in an interpretable CNN (as well as a few filters in traditional CNNs) may describe a specific object part without explicit names, e.g., the region of both the head and neck or the region connecting the torso and the tail. Therefore, in both [2] and [47], people did not require the inferred object region to describe the exact segment of a semantic part, and simply set a relatively loose criterion $\text{IoU}_{f,k}^I > 0.04$ or 0.2 to compute the filter interpretability.

TABLE 1
Part Interpretability of Filters in CNNs for Binary Classification of a Single Category Based on the VOC Part Dataset [4]

	bird	cat	cow	dog	horse	sheep	Avg.
AlexNet	0.332	0.363	0.340	0.374	0.308	0.373	0.348
AlexNet, interpretable	0.770	0.565	0.618	0.571	0.729	0.669	0.654
VGG-16	0.519	0.458	0.479	0.534	0.440	0.542	0.495
VGG-16, interpretable	0.818	0.653	0.683	0.900	0.795	0.772	0.770
VGG-M	0.357	0.365	0.347	0.368	0.331	0.373	0.357
VGG-M, interpretable	0.821	0.632	0.634	0.669	0.736	0.756	0.708
VGG-S	0.251	0.269	0.235	0.275	0.223	0.287	0.257
VGG-S, interpretable	0.526	0.366	0.291	0.432	0.478	0.251	0.390

Second, the location instability was proposed in [47]. The location instability of a filter is evaluated using the average deviation of distances between the inferred position and some ground-truth landmarks. There is also an assumption for this evaluation metric, i.e., the distance between an inferred part and a specific landmark should not change a lot through different images. As a result, people cannot set landmarks as the head and the tail of a snake, because the distance between different parts of a snake continuously changes when the snake moves.

Generally speaking, there are two advantages to use the location instability for evaluation:

- The computation of the location instability [47] is independent to the size of the receptive field (RF) of a neural activation. This solves a big problem with the evaluation of filter interpretability, i.e., state-of-the-art methods of computing a neural activation's image-resolution RFs (e.g., [50]) can only provide an approximate scale of the RF. The metric of location instability only uses central positions of part inferences of a filter, rather than use the entire inferred part segment, for evaluation. Thus, the location instability is a robust metric to evaluate the object-part interpretability of a filter.
- The location instability allows a filter to represent an object part without an explicit name (a half of the head). Nevertheless, the evaluation metric for filter interpretability is still an open problem.

TABLE 2
Part Interpretability of Filters in CNNs That are Trained for Multi-Category Classification Based on the VOC Part Dataset [4]

Network	Logistic log loss ³	Softmax log loss
VGG-16	0.710	0.723
interpretable	0.938	0.897
VGG-M	0.478	0.502
interpretable	0.770	0.734
VGG-S	0.479	0.435
interpretable	0.572	0.601

Filters in our interpretable CNNs exhibited significantly better part interpretability than ordinary CNNs in all comparisons.

TABLE 3
Location Instability of Filters ($\mathbb{E}_{f,k}[D_{f,k}]$) in CNNs That are Trained for the Binary Classification
of a Single Category Using the ILSVRC 2013 DET Animal-Part Dataset [43]

	gold.	bird	frog	turt.	liza.	koala	lobs.	dog	fox	cat	lion
AlexNet	0.161	0.167	0.152	0.153	0.175	0.128	0.123	0.144	0.143	0.148	0.137
AlexNet+ordinary layer	0.154	0.157	0.143	0.146	0.170	0.120	0.118	0.127	0.117	0.136	0.120
AlexNet, interpretable	0.084	0.095	0.090	0.107	0.097	0.079	0.077	0.093	0.087	0.095	0.084
VGG-16	0.153	0.156	0.144	0.150	0.170	0.127	0.126	0.143	0.137	0.148	0.139
VGG-16+ordinary layer	0.136	0.127	0.120	0.136	0.147	0.108	0.111	0.111	0.097	0.134	0.102
VGG-16, interpretable	0.076	0.099	0.086	0.115	0.113	0.070	0.084	0.077	0.069	0.086	0.067
VGG-M	0.161	0.166	0.151	0.153	0.176	0.128	0.125	0.145	0.145	0.150	0.140
VGG-M+ordinary layer	0.147	0.144	0.135	0.142	0.159	0.114	0.115	0.119	0.111	0.128	0.114
VGG-M, interpretable	0.088	0.088	0.089	0.108	0.099	0.080	0.074	0.090	0.082	0.103	0.079
VGG-S	0.158	0.166	0.149	0.151	0.173	0.127	0.124	0.143	0.142	0.148	0.138
VGG-S+ordinary layer	0.150	0.132	0.133	0.138	0.156	0.113	0.111	0.110	0.104	0.125	0.112
VGG-S, interpretable	0.087	0.101	0.093	0.107	0.096	0.084	0.078	0.091	0.082	0.101	0.082
	tiger	bear	rabb.	hams.	squi.	horse	zebra	swine	hippo.	catt.	sheep
AlexNet	0.142	0.144	0.148	0.128	0.149	0.152	0.154	0.141	0.141	0.144	0.155
AlexNet+ordinary layer	0.123	0.133	0.136	0.112	0.145	0.149	0.142	0.137	0.139	0.141	0.149
AlexNet, interpretable	0.090	0.095	0.095	0.077	0.095	0.098	0.084	0.091	0.089	0.097	0.101
VGG-16	0.144	0.143	0.146	0.125	0.150	0.150	0.153	0.141	0.140	0.140	0.150
VGG-16+ordinary layer	0.127	0.112	0.119	0.100	0.112	0.134	0.140	0.126	0.126	0.131	0.135
VGG-16, interpretable	0.097	0.081	0.079	0.066	0.065	0.106	0.077	0.094	0.083	0.102	0.097
VGG-M	0.145	0.144	0.150	0.128	0.150	0.151	0.158	0.140	0.140	0.143	0.155
VGG-M+ordinary layer	0.124	0.131	0.134	0.108	0.132	0.138	0.141	0.133	0.131	0.135	0.142
VGG-M, interpretable	0.089	0.101	0.097	0.082	0.095	0.095	0.080	0.095	0.084	0.092	0.094
VGG-S	0.142	0.143	0.148	0.128	0.146	0.149	0.155	0.139	0.140	0.141	0.155
VGG-S+ordinary layer	0.117	0.127	0.127	0.105	0.122	0.136	0.137	0.133	0.131	0.130	0.143
VGG-S, interpretable	0.089	0.097	0.091	0.076	0.098	0.096	0.080	0.092	0.088	0.094	0.101
	ante.	camel	otter	arma.	monk.	elep.	red pa.	gia.pa.			Avg.
AlexNet	0.147	0.153	0.159	0.160	0.139	0.125	0.140	0.125			0.146
AlexNet+ordinary layer	0.148	0.143	0.145	0.151	0.125	0.116	0.127	0.102			0.136
AlexNet, interpretable	0.085	0.102	0.104	0.095	0.090	0.085	0.084	0.073			0.091
VGG-16	0.144	0.149	0.154	0.163	0.136	0.129	0.143	0.125			0.144
VGG-16+ordinary layer	0.122	0.121	0.134	0.143	0.108	0.110	0.115	0.102			0.121
VGG-16, interpretable	0.091	0.105	0.093	0.100	0.074	0.084	0.067	0.063			0.085
VGG-M	0.146	0.154	0.160	0.161	0.140	0.126	0.142	0.127			0.147
VGG-M+ordinary layer	0.130	0.135	0.140	0.150	0.120	0.112	0.120	0.106			0.130
VGG-M, interpretable	0.077	0.104	0.102	0.093	0.086	0.087	0.089	0.068			0.090
VGG-S	0.143	0.154	0.158	0.157	0.140	0.125	0.139	0.125			0.145
VGG-S+ordinary layer	0.125	0.133	0.135	0.147	0.119	0.111	0.118	0.100			0.126
VGG-S, interpretable	0.077	0.102	0.105	0.094	0.090	0.086	0.078	0.072			0.090

Filters in our interpretable CNNs exhibited significantly lower localization instability than ordinary CNNs.

4.3.4 Robustness to Adversarial Attacks

In this experiment, we applied adversarial attacks [33] to both original CNNs and interpretable CNNs. The CNNs were learned to classify birds in the CUB200-2011 dataset and random images. Table 15 compares the average adversarial distortion $\sqrt{\frac{\sum_n (I'_i - I_i)^2}{n}}$ of the adversarial signal among all images between original CNNs and interpretable CNNs, where I represents the input image while I' denotes the adversarial counterpart. Because interpretable CNNs exclusively encoded object-part patterns and ignored textures, original CNNs usually exhibited stronger robustness to adversarial attacks than interpretable CNNs.

4.3.5 Experimental Results and Analysis

Feature interpretability of different CNNs is evaluated in Tables 1, 2, 3, 4, 5, and 6. Tables 1 and 2 show results based

on the metric in [2]. Tables 3, 4, and 5 list location instability of CNNs for binary classification of a single category. Table 6 reports location instability of CNNs that were learned for multi-category classification.

We compared our interpretable CNNs with two types of CNNs, i.e., the original CNN, the CNN with an additional conv-layer on the top (termed *AlexNet/VGG-16/VGG-M/VGG-S+ordinary layer*). To construct the CNN with a new conv-layer, we put a new conv-layer on the top of conv-layer. The filter size of the new conv-layer was $3 \times 3 \times \text{channel number}$, and output feature maps of the new conv-layer were in the same size of input feature maps. Because our interpretable CNN had an additional interpretable conv-layer, we designed the baseline CNN with a new conv-layer to enable fair comparisons. Our interpretable filters exhibited significantly higher part interpretability and lower location instability than traditional filters in baseline

TABLE 4
Location Instability of Filters ($\mathbb{E}_{f,k}[D_{f,k}]$) in CNNs That are Trained for Binary Classification of a Single Category Using the VOC Part Dataset [4]

	bird	cat	cow	dog	horse	sheep	Avg.
AlexNet	0.153	0.131	0.141	0.128	0.145	0.140	0.140
AlexNet+ordinary layer	0.147	0.125	0.139	0.112	0.146	0.143	0.136
AlexNet, interpretable w/o filter loss	0.091	0.090	0.091	0.089	0.086	0.088	0.089
AlexNet, interpretable	0.090	0.089	0.090	0.088	0.087	0.088	0.088
VGG-16	0.145	0.133	0.146	0.127	0.143	0.143	0.139
VGG-16+ordinary layer	0.125	0.121	0.137	0.102	0.131	0.137	0.125
VGG-16, interpretable w/o filter loss	0.099	0.087	0.102	0.078	0.096	0.101	0.094
VGG-16, interpretable	0.101	0.098	0.105	0.074	0.097	0.100	0.096
VGG-M	0.152	0.132	0.143	0.130	0.145	0.141	0.141
VGG-M+ordinary layer	0.142	0.120	0.139	0.115	0.141	0.142	0.133
VGG-M, interpretable w/o filter loss	0.089	0.095	0.091	0.086	0.086	0.091	0.090
VGG-M, interpretable	0.086	0.094	0.090	0.087	0.084	0.084	0.088
VGG-S	0.152	0.131	0.141	0.128	0.144	0.141	0.139
VGG-S+ordinary layer	0.137	0.115	0.133	0.107	0.133	0.138	0.127
VGG-S, interpretable w/o filter loss	0.093	0.096	0.093	0.086	0.083	0.090	0.090
VGG-S, interpretable	0.089	0.092	0.092	0.087	0.086	0.088	0.089

Filters in our interpretable CNNs exhibited significantly lower localization instability than ordinary CNNs.

CNNs over almost all comparisons. Table 7 reports the classification accuracy of different CNNs. Ordinary CNNs exhibited better performance in binary classification, while interpretable CNNs outperformed baseline CNNs in multi-category classification.

In addition, to prove the discrimination power of the learned filter, we further tested the average accuracy when we used the maximum activation score in a single filter's feature map as a metric for binary classification between birds in the CUB200-2011 dataset [36] and random images. In the scenario of classifying birds from random images, filters in the CNN was expected to learn the common appearance of birds, instead of summarizing knowledge from random images. Thus, we chose filters in the top conv-layer. If the maximum activation score of a filter exceeded a threshold, then we classified the input image as a bird; otherwise not. The threshold was set to the one that maximized the classification accuracy. Table 8 reports the average classification accuracy over all filters. Our interpretable filters outperformed ordinary filters.

Given a CNN for binary classification of an animal category in the VOC Part dataset [4], we manually annotated

TABLE 5
Location Instability of Filters ($\mathbb{E}_{f,k}[D_{f,k}]$) in CNNs for Binary Classification of a Single Category Using the CUB200-2011 Dataset

Neural network	Avg. location instability
AlexNet	0.150
AlexNet+ordinary layer	0.118
AlexNet, interpretable	0.070
VGG-16	0.137
VGG-16+ordinary layer	0.097
VGG-16, interpretable	0.076
VGG-M	0.148
VGG-M+ordinary layer	0.107
VGG-M, interpretable	0.065
VGG-S	0.148
VGG-S+ordinary layer	0.103
VGG-S, interpretable	0.073

the part name corresponding to the learned filters in the CNN. Table 9 reports the ratio of interpretable filters that corresponds to each object part.

Besides, we also analyzed samples that were incorrectly classified by the interpretable CNN. We used VGG-16 networks for the binary classification of an animal category in the VOC Part dataset [4]. We annotated the object-part name corresponding to each interpretable filter in the top interpretable layer. For each false positive sample without the target category, Fig. 9 localized the image regions that were incorrectly detected as specific object parts by interpretable filters. This figure helped people understand the reason for misclassification.

4.4 Effects of the Filter Loss

In this section, we evaluated effects of the filter loss. We compared the interpretable CNN learned with the filter loss with that without the filter loss (i.e., only using the mask layer without the filter loss).

TABLE 6
Location Instability of Filters ($\mathbb{E}_{f,k}[D_{f,k}]$) in CNNs That are Trained for Multi-Category Classification

	ILSVRC Part [43]	VOC Part [4]	
	Logistic log loss ³	Logistic log loss ³	Softmax log loss
VGG-16	–	0.128	0.142
ordinary layer	–	0.096	0.099
interpretable	–	0.073	0.075
VGG-M	0.167	0.135	0.137
ordinary layer	–	0.117	0.107
interpretable	0.096	0.083	0.087
VGG-S	0.131	0.138	0.138
ordinary layer	–	0.127	0.099
interpretable	0.083	0.078	0.082

Filters in our interpretable CNNs exhibited significantly lower localization instability than ordinary CNNs in all comparisons.

TABLE 7
Classification Accuracy Based on Different Datasets

Multi-category classification			
	ILSVRC Part	VOC Part	
	logistic ³	logistic ³	softmax
VGG-M	96.73	93.88	81.93
interpretable	97.99	96.19	88.03
VGG-S	96.98	94.05	78.15
interpretable	98.72	96.78	86.13
VGG-16	—	97.97	89.71
interpretable	—	98.50	91.60
Binary classification of a single category			
	ILSVRC Part	VOC Part	CUB200
AlexNet	96.28	95.40	95.59
interpretable w/o filter loss	—	93.98	—
interpretable	95.38	93.93	95.35
VGG-M	97.34	96.82	97.34
interpretable w/o filter loss	—	93.13	—
interpretable	95.77	94.17	96.03
VGG-S	97.62	97.74	97.24
interpretable w/o filter loss	—	93.83	—
interpretable	95.64	95.47	95.82
VGG-16	98.58	98.66	98.91
interpretable w/o filter loss	—	97.02	—
interpretable	96.67	95.39	96.51

In the binary classification of a single category, ordinary CNNs performed better, while in multi-category classification, interpretable CNNs exhibited superior performance.

TABLE 8
Classification Accuracy Based on a Single Filter

	filters in ordinary CNNs	filters in interpretable CNNs
AlexNet	68.7	75.1
VGG-M	69.9	80.2
VGG-16	72.1	82.4

We reported the average accuracy to demonstrate the discrimination power of individual filters.

4.4.1 Semantic Purity of Neural Activations

We proposed a metric to measure the semantic purity of neural activations of a filter. If a filter was activated at multiple locations besides the highest peak (i.e., the one corresponding to the target part), we considered this filter to have low semantic purity.

The semantic purity of a filter was measured as the ratio of neural activations within the range of the mask to all neural activations of the filter. In other words, a high semantic purity of a filter indicated that a filter mainly represented a single part, while a low semantic purity indicated that a filter represented multiple parts.

TABLE 9
Statistics of Semantic Meanings of Interpretable Filters

	bird	cow	cat	dog	horse	sheep
head	19.2	—	—	—	15.4	—
neck	21.2	—	—	5.8	—	—
torso	36.5	32.7	3.8	38.5	75.0	52.0
hip & tail	5.8	—	—	—	—	—
foot	5.8	—	—	—	9.6	3.8
wing	11.5	—	—	—	—	—
eye	—	30.8	55.8	11.5	7.7	—
nose & mouth	—	15.4	32.7	3.8	19.2	—
side face	—	9.6	—	—	—	—
leg	—	11.5	5.8	23.1	—	—
ear & horn	—	—	1.9	17.3	17.3	—

“—” indicates that the part is not selected as a label to describe the filter in a CNN. Except for CNNs for the bird and the horse, CNNs for other animals paid attention to detailed structures of the head. Thus, we annotated fine-grained parts inside the head for these CNNs.

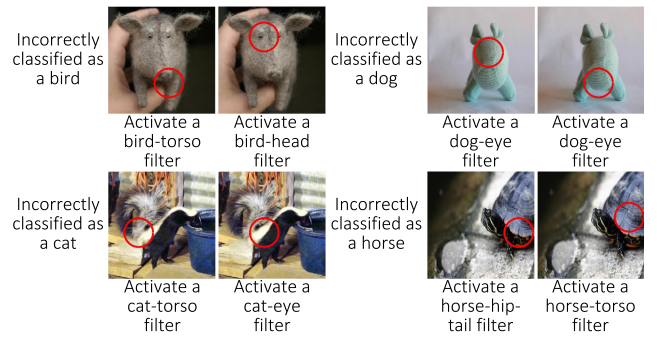


Fig. 9. Examples that were incorrectly classified by the interpretable CNN.

Let $x \in \mathbb{R}^{n \times n}$ be neural activations of a filter (after the ReLU layer and before the mask layer). The corresponding mask was given as $T_{\mu} \in \mathbb{R}^{n \times n}$. The semantic purity of neural activations was defined as $\text{purity} = \frac{\sum_j \sum_{i,j} \max(0, x_{ij}) \cdot \mathbf{1}(T_{\mu,ij} > 0)}{\sum_j \sum_{i,j} \max(0, x_{ij})}$.

$\mathbf{1}(\cdot)$ was the indicator function, which returns 1 if the condition in the braces was satisfied, and returns 0 otherwise. The semantic purity was supposed to be higher if neural activations were more concentrated.

We compared the semantic purity of neural activations between the interpretable CNN and the CNN trained with the mask layer but without filter loss. We constructed these CNNs based on the architectures of VGG-16, VGG-M and VGG-S, and learned the CNNs for binary classification on an animal category in the VOC Part dataset and the ILSVRC 2013 DET Animal-Part dataset. Experimental results are shown in Tables 10 and 11. It demonstrated that filters learned with the filter loss exhibited higher semantic purity

TABLE 10
Semantic Purity of Neural Activations of Interpretable Filters Learned With and Without the Filter Loss From the VOC Part Dataset

	bird	cat	cow	dog	horse	sheep	Avg.
VGG-16, interpretable	0.568	0.656	0.534	0.573	0.570	0.492	0.566
VGG-16 + mask layer, w/o filter loss	0.385	0.373	0.342	0.435	0.382	0.303	0.370
VGG-M, interpretable	0.444	0.616	0.395	0.540	0.408	0.387	0.465
VGG-M + mask layer, w/o filter loss	0.230	0.341	0.185	0.295	0.250	0.215	0.253
VGG-S, interpretable	0.418	0.437	0.390	0.398	0.421	0.369	0.406
VGG-S + mask layer, w/o filter loss	0.224	0.312	0.165	0.234	0.208	0.161	0.217

Filters learned with the filter loss exhibited significantly higher semantic purity than those learned without filter loss.

TABLE 11
Semantic Purity of Neural Activations of Interpretable Filters Learned the ILSVRC 2013 DET
Animal-Part Dataset With and Without the Filter Loss

VGG-16, interpretable	gold.	bird	frog	turt.	liza.	koala	lobs.	dog	fox	cat	lion
VGG-16 + mask layer, w/o filter loss	0.614	0.624	0.585	0.527	0.531	0.607	0.581	0.606	0.660	0.609	0.572
VGG-M, interpretable	0.331	0.441	0.342	0.352	0.362	0.339	0.344	0.482	0.449	0.393	0.344
VGG-M + mask layer, w/o filter loss	0.533	0.442	0.385	0.444	0.430	0.408	0.430	0.587	0.564	0.466	0.572
VGG-S, interpretable	0.307	0.251	0.249	0.198	0.266	0.255	0.212	0.384	0.393	0.389	0.293
VGG-S + mask layer, w/o filter loss	0.454	0.453	0.389	0.409	0.389	0.414	0.398	0.431	0.456	0.442	0.395
VGG-16, interpretable	0.245	0.265	0.186	0.189	0.216	0.212	0.209	0.316	0.330	0.284	0.197
VGG-16 + mask layer, w/o filter loss	0.583	0.614	0.535	0.679	0.677	0.559	0.532	0.530	0.564	0.466	0.572
VGG-M, interpretable	0.368	0.369	0.364	0.356	0.419	0.407	0.343	0.363	0.393	0.389	0.293
VGG-M + mask layer, w/o filter loss	0.455	0.442	0.472	0.452	0.434	0.397	0.421	0.412	0.420	0.425	0.420
VGG-S, interpretable	0.345	0.345	0.316	0.245	0.320	0.283	0.224	0.278	0.290	0.397	0.284
VGG-S + mask layer, w/o filter loss	0.448	0.419	0.411	0.405	0.416	0.430	0.469	0.384	0.403	0.439	0.404
VGG-16, interpretable	0.215	0.197	0.212	0.179	0.217	0.210	0.215	0.177	0.200	0.323	0.196
VGG-16 + mask layer, w/o filter loss	0.573	0.544	0.565	0.855	0.657	0.562	0.718	0.697			Avg.
VGG-M, interpretable	0.440	0.367	0.315	0.321	0.377	0.333	0.391	0.383			0.595
VGG-M + mask layer, w/o filter loss	0.584	0.435	0.441	0.419	0.400	0.399	0.541	0.468			0.459
VGG-S, interpretable	0.411	0.315	0.232	0.175	0.225	0.173	0.334	0.343			0.285
VGG-S + mask layer, w/o filter loss	0.471	0.386	0.386	0.436	0.394	0.408	0.459	0.471			0.420
VGG-16, interpretable	0.388	0.202	0.184	0.183	0.188	0.185	0.262	0.252			0.226

Filters learned with the filter loss exhibited significantly higher semantic purity than those learned without the filter loss.

than those learned without the filter loss. The filter loss forced each filter to exclusively represent a single object part during the training process.

4.4.2 Visualization of Filters

Besides the quantitative analysis of the semantic purity of neural activations, we visualized filter activations to compare the interpretable CNN and the CNN trained without the filter loss. Fig. 10 visualized neural activations of the filter in the first interpretable conv-layer of VGG-16 before the mask layer. Visualization results demonstrated that filters trained with the filter loss could generate more concentrated neural activations.

4.4.3 Location Instability

We compared the location instability among interpretable filters learned with the filter loss, those learned without the

filter loss, and ordinary filters. Here, we used filters in the first interpretable conv-layer of the interpretable CNN and filters in the corresponding conv-layer of the traditional CNNs for comparison.

We constructed the competing CNNs based on the VGG-16 architecture, and these CNNs were trained for single-category classification based on the VOC Part dataset and the ILSVRC 2013 DET Animal-Part dataset. As shown in Tables 12 and 13, the filter loss forced each filter to focus on a specific object part and reduced the location instability.

4.4.4 Activation Magnitudes

We further tested effects of the interpretable loss on neural activations among different categories. We used the VGG-M, VGG-S, and VGG-16 networks with either the logistic

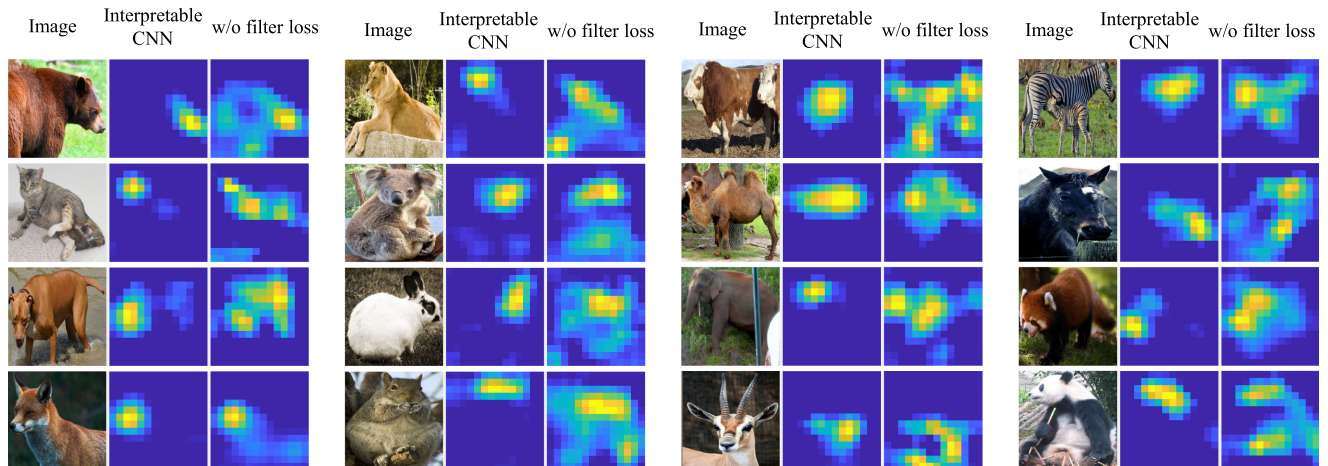


Fig. 10. Visualization of neural activations (before the mask layer) of interpretable filters learned with and without the filter loss. We visualized neural activations of the first interpretable conv-layer before the mask layer in the CNN. In comparison, visualization results in Fig. 5 correspond to feature maps after the mask layer. Filters trained with the filter loss tended to generate more concentrated neural activations and have higher semantic purity than filters learned without the filter loss.

TABLE 12

Location Instability of Filters ($\mathbb{E}_{f,k}[D_{f,k}]$) in the First Conv-Layer

	bird	cat	cow	dog	horse	sheep	Avg.
VGG-16	0.144	0.134	0.146	0.127	0.141	0.142	0.139
VGG-16 + mask layer, w/o filter loss	0.141	0.126	0.139	0.124	0.138	0.136	0.134
VGG-16, interpretable	0.130	0.127	0.134	0.109	0.131	0.126	0.126

The CNN was trained for the binary classification of a single category using the VOC Part dataset [4]. For the baseline, we added a conv-layer to the ordinary VGG-16 network, and selected the corresponding conv-layer in the network to enable fair comparisons. Interpretable filters learned with both the filter loss and the mask layer exhibited much lower localization instability than those learned with the mask layer but without the filter loss.

log loss or the softmax loss, which was trained to classify animal categories in the VOC Part dataset [4]. For each interpretable filter, given images of its target category, we recorded their neural activations (i.e., recording the maximal activation value in each of their feature maps). At the same time, we also recorded neural activations on other categories of the filter. The interpretable filter was supposed to activate much more strongly on its target category than on other (unrelated) categories. We collected all activation records on corresponding categories of all filters, and their mean value is reported in Table 14. In comparison, we also computed the mean value of all activations on unrelated categories of all filters in Table 14. This table shows that the interpretable filter was usually activated more strongly on the target category than on other categories. Furthermore, we also compared the proposed interpretable CNN with the ablation baseline *w/o filter loss*, in which the CNN was learned without the filter loss. Table 14 shows that the filter loss made each filter more prone to being triggered by a single category, i.e., boosting the feature interpretability.

5 CONCLUSION AND DISCUSSIONS

In this paper, we have proposed a general method to enhance feature interpretability of CNNs. We have designed a loss to push a filter in high conv-layers towards the representation of an object part during the learning process without any part annotations. Experiments have shown that each interpretable filter consistently represents a certain object

TABLE 14

The Mean Value of Neural Activations on the Target Categories and Those on Other Categories

Model	Logistic log loss		Softmax log loss	
	on target categories	on other categories	on target categories	on other categories
VGG-M, interpretable	107.9	7.5	6.2	1.1
w/o filter loss	18.6	6.3	4.7	1.1
VGG-S, interpretable	32.1	10.5	18.8	3.1
w/o filter loss	18.9	8.2	11.5	2.3
VGG-16, interpretable	948.9	81.5	106.7	5.1
w/o filter loss	40.5	12.9	67.0	6.0

Filters learned with the filter loss exhibited were usually more discriminative than those learned without the filter loss.

TABLE 15

Average Adversarial Distortion of the Original CNN and the Interpretable CNN

Model	Original CNN	Interpretable CNN
VGG-M	0.00302 \pm 0.00123	0.00243 \pm 0.00120
VGG-S	0.00305 \pm 0.00120	0.00266 \pm 0.00133
VGG-16	0.00293 \pm 0.00128	0.00280 \pm 0.00152

part of a category through different input images. In comparison, each filter in the traditional CNN usually represents a mixture of parts and textures.

Meanwhile, the interpretable CNN still has some drawbacks. First, in the scenario of multi-category classification, filters in a conv-layer are assigned with different categories. In this way, when we need to classify a large number of categories, theoretically, each category can only obtain a few filters, which will decrease a bit the classification performance. Otherwise, the interpretable conv-layer must contain lots of filters to enable the classification of a large number of categories. Second, the learning of the interpretable CNN has a strong assumption, i.e., each input image must contain a single object, which limits the applicability of the interpretable CNN. Third, the filter loss is only suitable to learn high conv-layers, because low conv-layers usually represent textures, instead of object parts. Finally, the interpretable CNN is not suitable to encode textural patterns.

TABLE 13

Location Instability of Filters ($\mathbb{E}_{f,k}[D_{f,k}]$) in the First Interpretable Conv-Layer

	gold.	bird	frog	turt.	liza.	koala	lobs.	dog	fox	cat	lion
VGG-16	0.152	0.152	0.142	0.150	0.167	0.125	0.126	0.139	0.137	0.149	0.134
VGG-16 + mask layer, w/o filter loss	0.143	0.148	0.141	0.146	0.162	0.121	0.120	0.137	0.130	0.142	0.130
VGG-16, interpretable	0.105	0.127	0.131	0.139	0.157	0.102	0.118	0.123	0.105	0.129	0.106
VGG-16	tiger	bear	rabb.	hams.	squi.	horse	zebra	swine	hippo.	catt.	sheep
VGG-16	0.143	0.139	0.143	0.124	0.145	0.147	0.155	0.138	0.139	0.142	0.148
VGG-16 + mask layer, w/o filter loss	0.137	0.131	0.140	0.119	0.134	0.143	0.142	0.133	0.131	0.134	0.145
VGG-16, interpretable	0.112	0.118	0.122	0.098	0.106	0.136	0.095	0.125	0.120	0.126	0.128
VGG-16	ante.	camel	otter	arma.	monk.	elep.	red pa.	gia.pa.			Avg.
VGG-16	0.141	0.144	0.155	0.156	0.134	0.127	0.139	0.126			0.142
VGG-16 + mask layer, w/o filter loss	0.134	0.142	0.147	0.153	0.129	0.124	0.125	0.116			0.136
VGG-16, interpretable	0.119	0.131	0.134	0.118	0.112	0.109	0.098	0.088			0.118

The CNN was trained for the binary classification of a single category using the ILSVRC 2013 DET Animal-Part dataset [43]. For the baseline, we added a conv-layer to the ordinary VGG-16 network, and selected the corresponding conv-layer in the network to enable fair comparisons. Interpretable filters learned with both the filter loss and the mask layer exhibited much lower localization instability than those learned with the mask layer but without the filter loss.

ACKNOWLEDGMENTS

This work was supported in part by National Natural Science Foundation of China (U19B2043 and 61906120), DARPA XAI Award N66001-17-2-4029, NSF IIS 1423305, and ARO project W911NF1810296.

REFERENCES

- [1] M. Aubry and B. C. Russell, "Understanding deep features with computer-generated imagery," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2015, pp. 2875–2883.
- [2] D. Bau, B. Zhou, A. Khosla, A. Oliva, and A. Torralba, "Network dissection: Quantifying interpretability of deep visual representations," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 3319–3327.
- [3] R. Chen, H. Chen, G. Huang, J. Ren, and Q. Zhang, "Explaining neural networks semantically and quantitatively," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2019, pp. 9186–9195.
- [4] X. Chen, R. Mottaghi, X. Liu, S. Fidler, R. Urtasun, and A. Yuille, "Detect what you can: Detecting and representing objects using holistic models and body parts," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2014, pp. 1979–1986.
- [5] X. Cheng, Z. Rao, Y. Chen, and Q. Zhang, "Explaining knowledge distillation by quantifying the knowledge," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2020.
- [6] A. Dosovitskiy and T. Brox, "Inverting visual representations with convolutional networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 4829–4837.
- [7] R. C. Fong and A. Vedaldi, "Interpretable explanations of black boxes by meaningful perturbation," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2017, pp. 3429–3437.
- [8] Y. Goyal, A. Mohapatra, D. Parikh, and D. Batra, "Towards transparent AI systems: Interpreting visual question answering models," 2016, *arXiv:1608.08974*.
- [9] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 770–778.
- [10] Z. Hu, X. Ma, Z. Liu, E. Hovy, and E. Xing, "Harnessing deep neural networks with logic rules," in *Proc. 54th Annu. Meeting Assoc. Comput. Linguistics*, vol. 1, 2016, pp. 2410–2420.
- [11] V. K. Ithapu, "Decoding the deep: Exploring class hierarchies of deep representations using multiresolution matrix factorization," in *Proc. CVPR Workshop Explainable Comput. Vis. Job Candidate Screening Competition*, 2017, pp. 1695–1704.
- [12] P. Koh and P. Liang, "Understanding black-box predictions via influence functions," in *Proc. 34th Int. Conf. Mach. Learn.*, 2017, pp. 1885–1894.
- [13] S. Kolouri, C. E. Martin, and H. Hoffmann, "Explaining distributed neural activations via unsupervised learning," in *Proc. CVPR Workshop Explainable Comput. Vis. Job Candidate Screening Competition*, 2017, pp. 1670–1678.
- [14] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," in *Proc. 25th Int. Conf. Neural Inf. Process. Syst.*, 2012, pp. 1097–1105.
- [15] D. Kumar, A. Wong, and G. W. Taylor, "Explaining the unexplained: A class-enhanced attentive response (CLEAR) approach to understanding deep neural networks," in *Proc. CVPR Workshop Explainable Comput. Vis. Job Candidate Screening Competition*, 2017, pp. 1686–1694.
- [16] H. Lakkaraju, E. Kamar, R. Caruana, and E. Horvitz, "Identifying unknown unknowns in the open world: Representations and policies for guided exploration," in *Proc. 31st AAAI Conf. Artif. Intell.*, 2017, pp. 2124–2132.
- [17] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proc. IEEE*, vol. 86, no. 11, pp. 2278–2324, Nov. 1998.
- [18] B. J. Lengerich, S. Konam, E. P. Xing, S. Rosenthal, and M. Veloso, "Visual explanations for convolutional neural networks via input resampling," in *Proc. ICML Workshop Vis. Deep Learn.*, 2017.
- [19] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," in *Proc. 31st Int. Conf. Neural Inf. Process. Syst.*, 2017, pp. 4768–4777.
- [20] A. Mahendran and A. Vedaldi, "Understanding deep image representations by inverting them," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2015, pp. 5188–5196.
- [21] C. Olah, A. Mordvintsev, and L. Schubert, "Feature visualization," *Distill*, 2017. [Online]. Available: <https://distill.pub/2017/feature-visualization>
- [22] M. T. Ribeiro, S. Singh, and C. Guestrin, "Why should i trust you?" explaining the predictions of any classifier," in *Proc. 22nd ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2016, pp. 1135–1144.
- [23] A. S. Ross, M. C. Hughes, and F. Doshi-Velez, "Right for the right reasons: Training differentiable models by constraining their explanations," in *Proc. 26th Int. Joint Conf. Artif. Intell.*, 2017, pp. 2662–2670.
- [24] S. Sabour, N. Frosst, and G. E. Hinton, "Dynamic routing between capsules," in *Proc. 31st Int. Conf. Neural Inf. Process. Syst.*, 2017, pp. 3859–3869.
- [25] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, "Grad-CAM: Visual explanations from deep networks via gradient-based localization," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2017, pp. 618–626.
- [26] Z. Si and S.-C. Zhu, "Learning AND-OR templates for object recognition and detection," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 35, no. 9, pp. 2189–2205, Sep. 2013.
- [27] M. Simon and E. Rodner, "Neural activation constellations: Unsupervised part model discovery with convolutional networks," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2015, pp. 1143–1151.
- [28] M. Simon, E. Rodner, and J. Denzler, "Part detector discovery in deep convolutional neural networks," in *Proc. Asian Conf. Comput. Vis.*, 2014, pp. 162–177.
- [29] K. Simonyan, A. Vedaldi, and A. Zisserman, "Deep inside convolutional networks: Visualising image classification models and saliency maps," 2013, *arXiv:1312.6034*.
- [30] K. Simonyan, A. Vedaldi, and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in *Proc. Int. Conf. Learn. Representations*, 2015.
- [31] S. Singh, A. Gupta, and A. A. Efros, "Unsupervised discovery of mid-level discriminative patches," in *Proc. Eur. Conf. Comput. Vis.*, 2012, pp. 73–86.
- [32] J. Su, D. V. Vargas, and S. Kouichi, "One pixel attack for fooling deep neural networks," *IEEE Trans. Evol. Comput.*, vol. 23, no. 5, pp. 828–841, 2019.
- [33] C. Szegedy et al., "Intriguing properties of neural networks," in *Proc. Int. Conf. Learn. Representations*, 2014.
- [34] J. Vaughan, A. Sudjianto, E. Brahimi, J. Chen, and V. N. Nair, "Explainable neural networks based on additive index models," 2018, *arXiv:1806.01933*.
- [35] C. Ventura, D. Masip, and A. Lapedriza, "Interpreting CNN models for apparent personality trait regression," in *Proc. CVPR Workshop Explainable Comput. Vis. Job Candidate Screening Competition*, 2017, pp. 1705–1713.
- [36] C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie, "The caltech-UCSD birds-200–2011 dataset," California Institute of Technology, Pasadena, CA, Tech. Rep. CNS-TR-2010-001, 2011.
- [37] A. S. Wicaksana and C. C. S. Liem, "Human-explainable features for job candidate screening prediction," in *Proc. CVPR Workshop Explainable Comput. Vis. Job Candidate Screening Competition*, 2017, pp. 1664–1669.
- [38] N. Wolchover and L. Reading, "New theory cracks open the black box of deep learning," *Quanta Mag.*, 2017. [Online]. Available: <https://www.quantamagazine.org/new-theory-cracks-open-the-black-box-of-deep-learning-20170921/>
- [39] J. Xie, W. Hu, S.-C. Zhu, and Y. N. Wu, "Learning inhomogeneous FRAME models for object patterns," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2014, pp. 1035–1042.
- [40] J. Yosinski, J. Clune, Y. Bengio, and H. Lipson, "How transferable are features in deep neural networks?" in *Proc. 27th Int. Conf. Neural Inf. Process. Syst.*, 2014, pp. 3320–3328.
- [41] M. D. Zeiler and R. Fergus, "Visualizing and understanding convolutional networks," in *Proc. Eur. Conf. Comput. Vis.*, 2014, pp. 818–833.
- [42] Q. Zhang, R. Cao, F. Shi, Y. Wu, and S.-C. Zhu, "Interpreting CNN knowledge via an explanatory graph," in *Proc. AAAI Conf. Artif. Intell.*, 2018, pp. 2124–2132.
- [43] Q. Zhang, R. Cao, Y. N. Wu, and S.-C. Zhu, "Growing interpretable part graphs on convnets via multi-shot learning," in *Proc. 31st AAAI Conf. Artif. Intell.*, 2016, pp. 2898–2906.
- [44] Q. Zhang, R. Cao, Y. N. Wu, and S.-C. Zhu, "Mining object parts from CNNs via active question-answering," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 3890–3899.


- [45] Q. Zhang, R. Cao, S. Zhang, M. Edmonds, Y. N. Wu, and S.-C. Zhu, "Interactively transferring CNN patterns for part localization," 2017, *arXiv: 1708.01783*.
- [46] Q. Zhang, W. Wang, and S.-C. Zhu, "Examining CNN representations with respect to dataset bias," in *Proc. AAAI Conf. Artif. Intell.*, 2018, pp. 4464–4473.
- [47] Q. Zhang, Y. N. Wu, and S.-C. Zhu, "Interpretable convolutional neural networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 8827–8836.
- [48] Q. Zhang, Y. Yang, H. Ma, and Y. N. Wu, "Interpreting CNNs via decision trees," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 6254–6263.
- [49] Q. Zhang and S.-C. Zhu, "Visual interpretability for deep learning: A survey," *Front. Inf. Technol. Electron. Eng.*, vol. 19, no. 1, pp. 27–39, 2018.
- [50] B. Zhou, A. Khosla, A. Lapedriza, A. Oliva, and A. Torralba, "Object detectors emerge in deep scene cnns," in *Proc. Int. Conf. Learn. Representations*, 2015.
- [51] B. Zhou, A. Khosla, A. Lapedriza, A. Oliva, and A. Torralba, "Learning deep features for discriminative localization," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 2921–2929.



Quanshi Zhang received the BS degree in machine intelligence from Peking University, China, in 2009, and the MS and PhD degrees in center for spatial information science from the University of Tokyo, Japan, in 2011 and 2014, respectively. He is currently an associate professor with the Shanghai Jiao Tong University, Shanghai, China. His research interests include computer vision and machine learning. In 2014, he went to the University of California, Los Angeles, Los Angeles, California, as a postdoctoral associate.



Xin Wang is currently working toward the PhD degree at the Shanghai Jiao Tong University, Shanghai, China. His main research interests include machine learning and computer vision.



Angeles, Los Angeles, California. His research interests include statistics, machine learning, and computer vision.



Huilin Zhou received the BS degree in mathematics from the University of Electronic Science and Technology of China, Chengdu, China, in 2019. Currently, she is working toward the PhD degree at the Shanghai Jiao Tong University, Shanghai, China. Her research interests include computer vision and machine learning.



Song-Chun Zhu received the PhD degree from Harvard University, Cambridge, Massachusetts, in 1996. He is currently a professor with the Departments of Statistics and Computer Science, University of California, Los Angeles, Los Angeles, California. He has published more than 300 papers in computer vision, statistical modeling and learning, cognition, language, robotics, and AI. He received a number of honors, including the Marr Prize, in 2003, the Aggarwal prize from the Int'l Association of Pattern Recognition, in 2008, the Holmoltz Test-of-Time prize, in 2013, twice Marr Prize honorary nominations, in 1999 and 2007, a Sloan Fellowship, the US NSF Career Award, and ONR Young Investigator Award, in 2001.

► For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/csdl.