

PICNN: A Pathway towards Interpretable Convolutional Neural Networks

Wengang Guo*, Jiayi Yang*, Huilin Yin, Qijun Chen, Wei Ye[†]

College of Electronic and Information Engineering, Tongji University, Shanghai, China
 {guowg, 2111125, yinhuilin, qjchen, yew}@tongji.edu.cn

Abstract

Convolutional Neural Networks (CNNs) have exhibited great performance in discriminative feature learning for complex visual tasks. Besides discrimination power, interpretability is another important yet under-explored property for CNNs. One difficulty in the CNN interpretability is that filters and image classes are entangled. In this paper, we introduce a novel pathway to alleviate the entanglement between filters and image classes. The proposed pathway groups the filters in a late conv-layer of CNN into class-specific clusters. Clusters and classes are in a one-to-one relationship. Specifically, we use the Bernoulli sampling to generate the filter-cluster assignment matrix from a learnable filter-class correspondence matrix. To enable end-to-end optimization, we develop a novel reparameterization trick for handling the non-differentiable Bernoulli sampling. We evaluate the effectiveness of our method on ten widely used network architectures (including nine CNNs and a ViT) and five benchmark datasets. Experimental results have demonstrated that our method PICNN (the combination of standard CNNs with our proposed pathway) exhibits greater interpretability than standard CNNs while achieving higher or comparable discrimination power.

Introduction

The remarkable discrimination power of convolutional neural networks (CNNs) fosters great applications in numerous tasks. However, in safety-critical domains such as autonomous vehicles (Zablocki et al. 2022) and healthcare (D’Amour et al. 2022), interpretability is another crucial property that needs to be considered.

The interpretability of CNNs has received growing interest in recent research. Earlier posthoc explanation methods (Zeiler and Fergus 2014; Simonyan, Vedaldi, and Zisserman 2013; Springenberg et al. 2014; Zhou et al. 2016; Selvaraju et al. 2017; Bau et al. 2017) focus on generating offline interpretations such as saliency maps (Simonyan, Vedaldi, and Zisserman 2013) and class activation mapping (CAM) (Zhou et al. 2016) for well-trained CNNs. But posthoc methods cannot improve the intrinsic interpretability of CNNs,

*These authors contributed equally.

[†]Corresponding author.

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

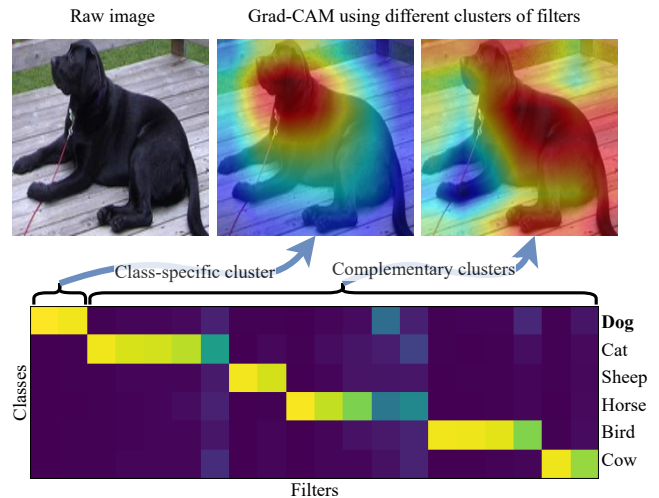


Figure 1: Top: comparison of Grad-CAM visualizations using different class-specific clusters of filters. Bottom: the display of the learned filter-class correspondence matrix \mathbf{P} . The dataset is PASCAL VOC Part (Chen et al. 2014) with six animal classes, and the target late conv-layer consists of 20 class-specific filters.

as they operate independently from the training process. Recently, the focus of the community has shifted to train interpretable CNNs. For example, (Zhang, Wu, and Zhu 2018; Shen et al. 2021) encourage each filter in a late conv-layer to respond to only one object part. Nevertheless, these methods cannot directly reveal concepts learned by filters, as they rely on some posthoc methods (Bau et al. 2017; Zhou et al. 2015) to associate filters to the predefined concepts.

In this paper, interpretability is characterized by the alignment between filters and class-level concepts that are human-understandable. We propose a novel method to train interpretable CNNs whose filters in the late conv-layer can directly reveal class-level concepts without the help of any posthoc methods. We focus on filters in the late conv-layer since these filters encode class-level concepts rather than low-level primitives (Zeiler and Fergus 2014). (Fong and Vedaldi 2018) has shown that a CNN uses multiple filters in the late conv-layers to collectively encode a class-level

concept. Considering a target late conv-layer with N filters, there exist 2^N possible combinations of filters, allowing these filters to encode up to 2^N class-level concepts. This huge concept space prohibits the explanation of the relations between filters and classes due to their complex many-to-many correspondences, known as filter-class entanglement. As argued in prior work (Liang et al. 2020), filter-class entanglement is one of the most critical reasons that hamper the interpretability of CNNs.

To mitigate filter-class entanglement, we group the N filters into K class-specific clusters, where K is the number of image classes. Clusters and classes are in a one-to-one correspondence relation, i.e., each class-specific cluster is encouraged to contain filters that encode the class-level concepts of only one particular image class. A direct way to group filters is to predefine a fixed filter-cluster assignment matrix before training. However, this results in suboptimal performance, as the optimal number of filters assigned to each class is unknown a priori. (Shen et al. 2021) employs spectral clustering (Shi and Malik 2000) to group filters during each forward propagation. Differing from our paper, their goal is to make each filter represent a set of interpretable image features, such as a specific object part or image region with a clear meaning.

In this paper, we develop a novel pathway that can group filters into class-specific clusters and be applied to many network architectures. Specifically, the pathway introduces a learnable matrix $\mathbf{P} \in \mathbb{R}^{K \times N}$ whose element $p_{y,i}$ represents the probabilistic correspondence between the i -th filter and y -th class. Based on \mathbf{P} , we can use the argmax operation to generate a binary filter-cluster assignment matrix. However, the argmax operation is non-differentiable, impeding end-to-end optimization. One potential solution is using Gumbel-Softmax (Jang, Gu, and Poole 2016) to approximate the argmax operation to enable gradient propagation. Nevertheless, this approximation introduces bias, as Gumbel-Softmax constitutes a biased estimate of the original argmax operation. Such bias can accumulate during the training process and result in suboptimal performance. To address this challenge, we consider $p_{y,j}$ as the parameter of the Bernoulli distribution, based on which a filter-cluster assignment matrix is sampled from \mathbf{P} . Since the sampling is also non-differentiable, we propose a novel reparameterization trick to support end-to-end training. Our filter-cluster assignment strategy yields superior performance compared to the Gumbel-Softmax-based strategy. During end-to-end training, filter grouping and filter learning are jointly optimized: Filter grouping is conducted in the forward propagation, while filter learning is conducted in the backward propagation. Figure 1 displays the learned correspondence matrix \mathbf{P} and Grad-CAM (Selvaraju et al. 2017) visualizations for validating the learned filter-class correspondence.

The contributions of this paper are as follows:

- We develop a novel pathway to transform a standard CNN into an interpretable CNN. This pathway is versatile and can be flexibly combined with many CNN architectures and even Transformer architectures.
- We propose to group filters into class-specific clusters

by the filter-cluster assignment matrix, which is sampled from a learnable filter-class correspondence matrix \mathbf{P} by the Bernoulli distribution. Each element of \mathbf{P} can be used as the parameter of the Bernoulli distribution.

- We propose a reparameterization trick for sampling by the Bernoulli distribution to support end-to-end training.
- We evaluate the effectiveness of our proposed PICNN using five benchmark datasets. Experimental results demonstrate that PICNN achieves higher or comparable discrimination power and better interpretability than backbone models.

Related Work

We briefly survey the literature on posthoc filter interpretability, learning interpretable filters, and class-specific filters.

Posthoc filter interpretability has been widely studied, which aims to build the mapping between the abstract patterns of filters in well-trained CNNs and the human-understandable domains such as images (Wang et al. 2021). Earlier works (Zeiler and Fergus 2014; Mahendran and Vedaldi 2015; Yosinski et al. 2015) inspect maximal activations of filters across different images. CAM (Zhou et al. 2016) localizes image regions that are most important for a target class. CAM has gained much popularity and promotes numerous further studies (Selvaraju et al. 2017; Lee et al. 2021; Hasany, Petitjean, and Mériaudeau 2023; Sarkar et al. 2023). Notably, Grad-CAM (Selvaraju et al. 2017) extends CAM by using gradients to weigh the contribution of each filter to the target class. Other well-known posthoc methods include guided backpropagation (Springenberg et al. 2014), saliency maps (Simonyan, Vedaldi, and Zisserman 2013), and deconvolutional network (Zeiler and Fergus 2014; Dosovitskiy and Brox 2016). Moreover, some works disentangle the filter representations of a well-trained CNN into an explanatory graph (Zhang et al. 2018), a decision tree (Zhang et al. 2019), or textual descriptions (Hendricks et al. 2016; Yang, Kim, and Joo 2022). However, these posthoc methods cannot remove the existing filter-class entanglement in well-trained CNNs and may not faithfully capture what the original CNNs compute (Rudin 2019). In contrast, our work focuses on training interpretable CNNs, where the relations between filters and classes are clearly revealed by the learned filter-class correspondence matrix.

Learning interpretable filters has been conducted in ICNN (Zhang, Wu, and Zhu 2018). The learned each filter represents a specific object part, such as animal eyes. Filters in ICNN could only represent object parts in ball-like areas. To overcome this limitation, ICCNN (Shen et al. 2021) extends filter interpretability to image regions with arbitrary shapes. However, the filters learned in both ICNN and ICCNN are not class-specific, as they are trained in a class-agnostic manner. Learning class-specific filters is first introduced by Class-Specific Gate (CSG) (Liang et al. 2020). Specifically, CSG introduces a binary filter-class gate matrix to represent correspondences between filters and classes. The binary filter-class gate matrix is relaxed by the l_1 norm

to weigh feature maps of filters and simultaneously forced to be sparse. Differing from CSG, we consider the filter-class correspondences as probability parameters and optimize these parameters using a probability approach.

Class-specific filters have been widely utilized across various studies. Class-specific filters enable the integration of direct supervision into the hidden layers of CNNs, thereby alleviating gradient vanishing (Jiang et al. 2017). (Wang, Morariu, and Davis 2018; Martinez et al. 2019) improve fine-grained recognition using a filter bank that captures class-specific discriminative patches. Except for discriminative models, generative models also benefit from class-specific formulation. (Tang et al. 2020; Li et al. 2021) create multiple class-specific generators within a GAN (Goodfellow et al. 2014) to facilitate the generation of small objects. (Kweon et al. 2021) proposes a class-specific adversarial erasing framework to generate more precise CAM. Unlike these models that enhance the discrimination or generation power of CNNs, we focus on filter interpretability. Additionally, these models require predefined filter-class correspondences, whereas our work automatically learns filter-class correspondences.

Model PICNN

This paper aims to train interpretable CNNs whose filters are disentangled from image classes and can directly reveal class-level concepts. To achieve this, we propose to group filters into class-specific clusters. Each class-specific cluster of filters exclusively activates for inputs from a particular class and deactivates for inputs from other classes.

Objective Function

We enrich the target conv-layer in a CNN, typically the last conv-layer, with an extra pathway, i.e., the interpretation pathway to learn class-specific filters (Figure 2). We preserve the original discrimination pathway to fit the underlying task.

1) In the discrimination pathway (the blue shaded area in Figure 2), the classifier (i.e., the multilayer perceptron (MLP)) receives complete feature maps from the target conv-layer and outputs a softmax-normalized prediction. The objective is to fit the underlying task by optimizing a loss function, such as the cross-entropy loss for a classification task:

$$\mathcal{L}_{\text{dis}} = H(\mathbf{y}, \mathbf{y}_1) \quad (1)$$

where $\mathbf{y} \in \mathbb{R}^K$ is the one-hot encoding of the class label¹, $\mathbf{y}_1 \in \mathbb{R}^K$ is the prediction from the discrimination pathway, and K is the number of image classes.

2) In the interpretation pathway (the green shaded area in Figure 2), we input into the same classifier used in the discrimination pathway the feature maps of a class-specific cluster of filters. We postulate that each class-specific cluster of filters specialized for a given class is essential for the network to classify that class, whereas other clusters can

be removed/masked without compromising discrimination power. Then, the classifier outputs a prediction $\mathbf{y}_2 \in \mathbb{R}^K$. The cross-entropy loss in this interpretation pathway is:

$$\mathcal{L}_{\text{int}} = H(\mathbf{y}, \mathbf{y}_2) \quad (2)$$

The interpretation pathway is optimized to achieve a comparable classification performance as the discrimination pathway. The optimization of the interpretation pathway encourages the learning of the class-specific filters in the discrimination pathway. We simultaneously optimize the losses of these two pathways as follows:

$$\mathcal{L} = \mathcal{L}_{\text{dis}} + \lambda \mathcal{L}_{\text{int}} \quad (3)$$

where λ is a regularization parameter.

Class-specific Grouping of Filters

Filter-class Correspondence Matrix We introduce a learnable matrix $\mathbf{P} \in \mathbb{R}^{K \times N}$ to indicate the probabilities of the filter-class correspondences, where N is the number of filters in the target conv-layer. A larger value of $p_{y,i}$ represents a stronger correspondence between the i -th filter and the y -th class. With the correspondence matrix \mathbf{P} , we can generate the filter-cluster assignment matrix \mathbf{Z} by the argmax operation. However, this operation is non-differentiable. Using Gumbel-Softmax (Jang, Gu, and Poole 2016) for approximating the argmax operation induces approximation bias and yields suboptimal performance. Thus, we propose to sample \mathbf{Z} from \mathbf{P} by the Bernoulli distribution parameterized by the element of \mathbf{P} .

Bernoulli Sampling Supposing the input belongs to the y -th class, we conduct N independent Bernoulli trials to generate a binary filter-cluster assignment vector $\mathbf{z} = [z_1, \dots, z_N] \in \{0, 1\}^N$ from $\mathbf{p}_y = [p_{y,1}, \dots, p_{y,N}]$, i.e., $z_i \sim \text{Ber}(p_{y,i}), i = 1, 2, \dots, N$, where \mathbf{p}_y is the y -th row of matrix \mathbf{P} . The i -th element z_i determines whether the i -th filter belongs to (if $z_i = 1$) the class-specific cluster of the y -th class or not (if $z_i = 0$). After that, we use \mathbf{z} to mask the complete feature maps $\mathbf{H} = \{\mathbf{H}_i\}_{i=1}^N$ (each \mathbf{H}_i is a 2D matrix) of all the filters to generate masked feature maps $\tilde{\mathbf{H}} = \{\tilde{\mathbf{H}}_i\}_{i=1}^N$ by the Hadamard product:

$$\tilde{\mathbf{H}}_i = \mathbf{H}_i \odot \mathbf{z}_i \quad i = 1, \dots, N \quad (4)$$

where \mathbf{z}_i is broadcasted along the spatial dimensions of \mathbf{H}_i for compatible product operation. If $z_i = 1$, $\tilde{\mathbf{H}}_i$ is the original feature maps of the i -th filter; If $z_i = 0$, $\tilde{\mathbf{H}}_i$ is set to a matrix of all zero elements. Finally, we feed both the complete feature maps \mathbf{H} and the masked ones $\tilde{\mathbf{H}}$ into the classifier to compute predictions \mathbf{y}_1 and \mathbf{y}_2 , respectively.

Optimization

Now we elaborate on how to optimize the correspondence matrix \mathbf{P} together with other network parameters using stochastic gradient descent (SGD). The challenge is that sampling the filter-cluster assignment matrix \mathbf{Z} from \mathbf{P} by the Bernoulli trials is non-differentiable, which impedes SGD optimization of \mathbf{P} . To deal with this, we resort to the reparameterization trick (Kingma and Welling 2014).

¹In the paper, integer $y \in \{1, \dots, K\}$ represents the true class label of an image, while vector $\mathbf{y} \in \mathbb{R}^K$ represents the one-hot encoding of y .

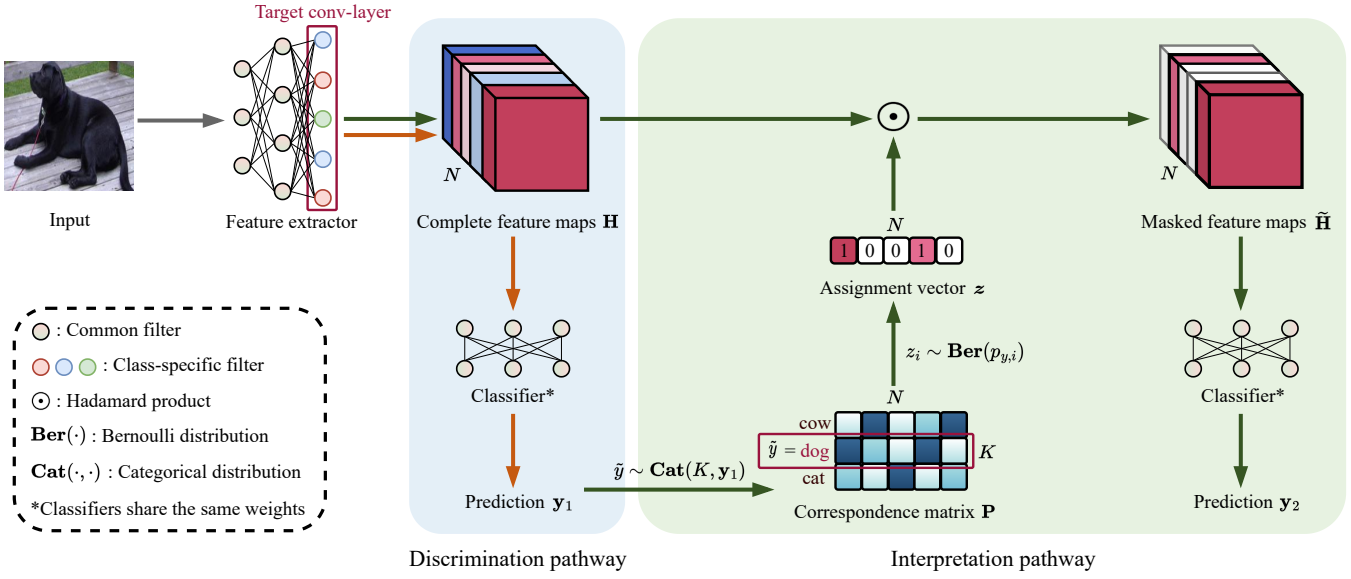


Figure 2: The pipeline of our model PICNN, which consists of (a) the discrimination pathway to fit the underlying classification task and (b) the interpretation pathway to group filters into class-specific clusters, whose optimization encourages the discrimination pathway to learn class-specific filters.

To avoid computing the partial derivative for sampling, we propose to reparameterize the assignment vector \mathbf{z} , which is initially drawn from the Bernoulli distribution with the probability parameter \mathbf{p}_y . Specifically, we redefine $z_i \sim \text{Ber}(p_{y,i})$ as:

$$z_i = p_{y,i} + \eta_i \quad (5)$$

where η_i is a normalization scalar defined as:

$$\eta_i = \begin{cases} 1 - p_{y,i}, & p_{y,i} \geq \epsilon_i \\ -p_{y,i}, & p_{y,i} < \epsilon_i \end{cases} \quad i = 1, \dots, N \quad (6)$$

where ϵ_i is sampled from the continuous uniform distribution $\mathcal{U}(0, 1)$, i.e., $\epsilon_i \sim \mathcal{U}(0, 1)$. This reparameterization trick converts the sampling of z_i from a Bernoulli distribution, which relies on trainable parameters, into sampling ϵ_i from a fixed uniform distribution. This trick enables us to compute the partial derivative $\frac{\partial \mathbf{z}}{\partial \mathbf{P}}$.

Theorem 1. $z_i \sim \text{Ber}(p_{y,i})$.

Proof. Substituting Equation 6 into Equation 5, we obtain:

$$z_i = \begin{cases} 1, & p_{y,i} \geq \epsilon_i \\ 0, & p_{y,i} < \epsilon_i \end{cases} \quad i = 1, \dots, N \quad (7)$$

Since $\epsilon_i \sim \mathcal{U}(0, 1)$, the probability $p(p_{y,i} \geq \epsilon_i)$ equals the ratio between the length of the interval $[0, p_{y,i}]$ and the total interval $[0, 1]$. Thus, we have $p(z_i = 1) = p(p_{y,i} \geq \epsilon_i) = p_{y,i}$ and $p(z_i = 0) = 1 - p(z_i = 1) = 1 - p_{y,i}$, meaning z_i follows the Bernoulli distribution $\text{Ber}(p_{y,i})$. \square

Our model may suffer from a trivial solution, as the prediction \mathbf{y}_2 is indirectly dependent on the label y (from which a red arrow flows in Figure 3). As a result, the network can exploit this dependency and effortlessly minimize the loss \mathcal{L}_{int} (Figure 4(a)). To prevent this, we propose to replace the

label y with a pseudo-label \tilde{y} for indexing the correspondence matrix \mathbf{P} . The pseudo-label \tilde{y} is drawn from a categorical distribution parameterized by the softmax-normalized prediction \mathbf{y}_1 from the discrimination pathway, i.e., $\tilde{y} \sim \text{Cat}(K, \mathbf{y}_1)$. This intuitive idea is that the pseudo-label \tilde{y} gradually approaches the label y as the training unfolds. To support SGD optimization, we utilize the reparameterization trick once again to redefine \tilde{y} . Specifically, we reparameterize the pseudo-label as:

$$\tilde{y} = \mathbf{y}_1 + \boldsymbol{\tau} \quad (8)$$

where $\boldsymbol{\tau}$ is a constant vector defined as:

$$\boldsymbol{\tau}[k] = \begin{cases} 1 - \mathbf{y}_1[k], & \sum_{j=0}^{k-1} \mathbf{y}_1[j] < \xi \leq \sum_{j=0}^k \mathbf{y}_1[j] \\ -\mathbf{y}_1[k], & \text{otherwise} \end{cases} \quad k = 1, \dots, K \quad (9)$$

where $\xi \sim \mathcal{U}(0, 1)$, $[k]$ indexes the k -th entry of a vector and we append 0 before the first element of \mathbf{y}_1 and 1 after the last element of \mathbf{y}_1 .

Theorem 2. $\tilde{y} \sim \text{Cat}(K, \mathbf{y}_1)$.

Proof. Substituting Equation 9 into Equation 8, we obtain:

$$\tilde{y}[k] = \begin{cases} 1, & \sum_{j=0}^{k-1} \mathbf{y}_1[j] < \xi \leq \sum_{j=0}^k \mathbf{y}_1[j] \\ 0, & \text{otherwise} \end{cases} \quad (10)$$

Since $\xi \sim \mathcal{U}(0, 1)$, the probability $p(\sum_{i=0}^{k-1} \mathbf{y}_1[i] < \xi \leq \sum_{i=0}^k \mathbf{y}_1[i])$ equals to the ratio between the length of the interval $[\sum_{i=0}^{k-1} \mathbf{y}_1[i], \sum_{i=0}^k \mathbf{y}_1[i]]$ and the total interval $[0, 1]$. Thus, we have $p(\tilde{y} = k) = \mathbf{y}_1[k]$, i.e., $\tilde{y} \sim \text{Cat}(K, \mathbf{y}_1)$. \square

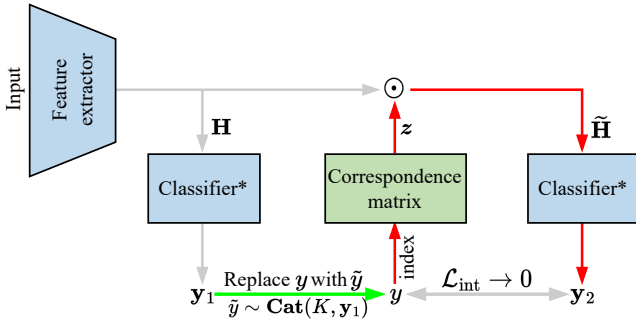


Figure 3: The interpretation pathway encounters a label leakage issue, as the label y is employed to index the correspondence matrix \mathbf{P} . To deal with this issue, we replace y with \tilde{y} sampled from the categorical distribution parameterized by the softmax-normalized prediction of the y -th class in the discrimination pathway.




Figure 4: Trivial solution. (a) Using the label y as index vector, the loss \mathcal{L}_{int} quickly becomes zero, leading to a trivial solution. (b) After replacing the label y with the pseudo-label \tilde{y} , our model shows a stable evolution.

Figure 4(b) demonstrates that this categorical sampling effectively circumvents the trivial solution.

Time Complexity Analysis

Most clustering methods incur a high time complexity between $\mathcal{O}(M^2)$ and $\mathcal{O}(M^3)$ if the number of input images is M . In PICNN, the interpretation pathway plays a role that is equivalent to a clustering component but with a lower time complexity of $\mathcal{O}(M(2K + 2K^2 + 2KN + Nd^2 + 2KN^2))$, where K is the number of image classes, N is the number of filters and d is the size of the feature map. And the space complexity is $\mathcal{O}(NK)$.

Experiments

Experimental Settings

Evaluation Metric We use classification accuracy and mutual information score (MIS) (Liang et al. 2020) as evaluation metrics to assess the discrimination power and interpretability. We report the following four metrics: **1)** classification accuracy of the discrimination pathway ($\text{ACC1}\uparrow$). ACC1 evaluates the discrimination power. **2)** classification accuracy of the interpretation pathway ($\text{ACC2}\uparrow$). Since the input to the interpretation pathway is the feature maps of a class-specific cluster of filters, ACC2 is consistently lower

than ACC1 . **3)** classification accuracy using all the filters except a class-specific cluster of filters ($\text{ACC3}\downarrow$). Since the filters used in the computation of ACC3 and ACC2 complement each other, the lower ACC3 is, the better the model. **4)** MIS ($\text{MIS}\uparrow$) measures the mutual information between filter activation and prediction on classes. MIS is defined as $\text{MIS} = \text{mean}_i(\max_y(m_{yi}))$, where $m_{yi} = \text{MI}(\mathbf{H}_i\|\mathbf{y})$. The metrics ACC2 , ACC3 , and MIS quantify the degree of filter-class entanglement from different aspects, thus they can evaluate the interpretability.

Datasets and Types of Backbone CNNs We use three benchmark classification datasets in Table 1, including CIFAR-10 (Krizhevsky, Hinton et al. 2009), STL-10 (Coates, Ng, and Lee 2011), and PASCAL VOC Part (Chen et al. 2014). CIFAR-10 consists of 50,000 training images and 10,000 test images in 10 classes. STL-10 contains 5,000 training images and 8,000 test images in 10 classes. Following (Liang et al. 2020), we select six animal classes from PASCAL VOC Part with a 70%/30% training/test split. To further evaluate the efficacy and effectiveness of PICNN on large datasets with more classes, we use CIFAR-100 (Krizhevsky, Hinton et al. 2009) and TinyImageNet (Deng et al. 2009) in Table 2. Like CIFAR-10, CIFAR-100 also consists of 50,000 training images and 10,000 test images evenly distributed into 100 classes. TinyImageNet is a scaled-down version of the original ImageNet involving 200 classes with a 10:1 ratio of training images to test images. We use the official training/test data split, except for PASCAL VOC Part.

We combine our interpretation pathway with six typical CNN architectures, three large CNN architectures, and a transformer architecture (Dosovitskiy et al. 2020) (ViT-b-12) to make them interpretable. The six CNN architectures are VGG-11 (Simonyan and Zisserman 2014), AlexNet (Krizhevsky 2014), ResNet-18 (He et al. 2016), DenseNet-121 (Huang et al. 2017), MobileNetV2 (Sandler et al. 2018), and EfficientNet-B0 (Tan and Le 2019). The three deeper and wider CNN architectures include ResNet-50 (He et al. 2016), ResNet-152 (He et al. 2016), and Wide-ResNet (Zerhouni et al. 2017).

Implementation Details Our code is based on the PyTorch (Paszke et al. 2019) toolbox and publicly available at Github². We make filters in the target conv-layer class-specific by optimizing our loss function. The regularization parameter λ is set to 2 and the effect of the λ values is discussed later. Other default settings include: a batch size of 128; the Adam optimizer with an initial learning rate of 0.001; pretrained weights from ImageNet (Deng et al. 2009); and a total of 200 training epochs. All metrics presented in this paper are computed on the test set. The experiments are carried out on a server with an Xeon(R) Platinum 8352V CPU and one Nvidia RTX 4090 GPU.

Experimental Results

Comparison with Existing Methods We first compare our PICNN with a Standard CNN (STD) (ResNet-18, $\lambda =$

²<https://github.com/spdj2271/PICNN>

		CIFAR-10				PASCAL VOC Part				STL-10			
Backbone	Model	ACC1↑	ACC2↑	ACC3↓	MIS↑	ACC1↑	ACC2↑	ACC3↓	MIS↑	ACC1↑	ACC2↑	ACC3↓	MIS↑
ResNet-18	STD	0.952	0.588	0.950	0.191	0.782	0.769	0.781	0.203	0.895	0.778	0.894	0.148
	ICCNN	0.955	N/A	N/A	0.067	0.754	N/A	N/A	0.140	0.871	N/A	N/A	0.052
	CSG	0.830	0.505	0.659	0.132	0.781	0.686	0.666	0.148	0.888	0.373	0.699	0.115
	PICNN	0.951	0.948	0.333	0.243	0.800	0.798	0.639	0.246	0.896	0.893	0.657	0.224
VGG-11	STD	0.918	0.401	0.913	0.128	0.803	0.629	0.799	0.138	0.702	0.684	0.697	0.142
	PICNN	0.915	0.910	0.192	0.177	0.840	0.837	0.346	0.166	0.716	0.709	0.316	0.131
AlexNet	STD	0.904	0.308	0.893	0.091	0.838	0.733	0.837	0.122	0.646	0.618	0.644	0.115
	PICNN	0.905	0.900	0.188	0.191	0.821	0.817	0.432	0.169	0.645	0.641	0.205	0.118
DenseNet-121	STD	0.958	0.438	0.954	0.163	0.918	0.917	0.714	0.148	0.829	0.791	0.824	0.200
	PICNN	0.957	0.957	0.462	0.257	0.917	0.915	0.763	0.229	0.832	0.831	0.734	0.266
MobileNet V2	STD	0.950	0.540	0.948	0.208	0.930	0.844	0.928	0.176	0.829	0.818	0.828	0.210
	PICNN	0.956	0.952	0.403	0.261	0.928	0.925	0.844	0.251	0.839	0.834	0.746	0.258
EfficientNet-B0	STD	0.955	0.351	0.951	0.174	0.942	0.889	0.940	0.165	0.836	0.820	0.828	0.229
	PICNN	0.956	0.955	0.380	0.253	0.940	0.937	0.873	0.253	0.842	0.833	0.764	0.275

Table 1: Experimental results on three benchmark datasets. We conduct comparison experiments between standard methods and our method on six CNN architectures. STD stands for standard CNNs, ICCNN stands for Compositional CNN, CSG stands for Class-Specific Gate CNN, and PICNN is our method. A higher value is better for ACC1, ACC2, and MIS, while a lower value is better for ACC3.

Dataset	Class	Model	ACC1↑	ACC2↑	ACC3↓	MIS↑
CIFAR-100	100	STD	0.755	0.039	0.753	0.020
		PICNN	0.724	0.707	0.131	0.026
TinyImageNet	200	STD	0.619	0.010	0.619	0.010
		PICNN	0.555	0.514	0.082	0.017

Table 2: Performance on large multi-class datasets with ResNet-18 as the backbone.

0), Interpretable Compositional CNNs (ICCNN) (Shen et al. 2021) and Class-Specific Gate (CSG) (Liang et al. 2020). All the networks use the same backbone ResNet-18. For STD, we randomly initialize the correspondence matrix \mathbf{P} and use our Bernoulli sampling to generate the filter-cluster assignment matrix. As shown in the upper part of Table 1, PICNN significantly outperforms the comparison methods in terms of ACC2, ACC3, and MIS across all datasets. Since the filters learned by ICCNN are not class-specific, ACC2 and ACC3 are not suitable to evaluate its performance. The corresponding result is denoted as N/A. Additionally, PICNN is better than STD on PASCAL VOC Part and STL-10 datasets in terms of ACC1 and achieves comparable results to STD on CIFAR-10. CSG shows a noticeable decrease in ACC1 on CIFAR-10, and ICCNN shows relatively poorer performance on PASCAL VOC Part and STL-10. The results on large datasets are shown in Table 2 and demonstrate once again that PICNN has a significant improvement in ACC2, ACC3, and MIS over STD.

All these results indicate that PICNN achieves better interpretability and maintains comparable discrimination

Backbone	Model	ACC1↑	ACC2↑	ACC3↓	MIS↑
ResNet-50	STD	0.951	0.457	0.948	0.115
	PICNN	0.955	0.950	0.318	0.231
ResNet-152	STD	0.957	0.481	0.954	0.150
	PICNN	0.958	0.954	0.309	0.230
Wide-ResNet	STD	0.940	0.295	0.930	0.117
	PICNN	0.940	0.933	0.137	0.204
ViT-b-12	STD	0.799	0.242	0.794	0.066
	PICNN	0.802	0.746	0.039	0.096

Table 3: Evaluation for large CNNs and Vision Transformer architecture on the CIFAR-10 dataset.

power simultaneously. In contrast, the comparison methods are unsatisfactory in both the discrimination power and interpretability.

Various Network Architectures One advantage of our method lies in its versatility to combine with various network architectures. As shown in Table 1 and Table 3, our novel pathway significantly boosts the interpretability of all the ten network architectures, which include six typical CNN architectures, three large CNN architectures, and a transformer architecture. Except for using DenseNet-121 as the backbone on the PASCAL VOC Part dataset, PICNN has the best performance in terms of all three interpretability metrics ACC2, ACC3, and MIS. Compared with STD, the improvement is significant. PICNN does not work well when DenseNet-121 is the backbone. One reason might be:

Backbone	Model	Train	Infer	Model	Train	Infer
ResNet-50	STD	18302	1564	PICNN	18484	1592
ResNet-152	STD	27339	2403	PICNN	27377	2514
Wide-ResNet	STD	76572	4733	PICNN	77212	4831
ViT-b-12	STD	37245	3234	PICNN	37615	3257

Table 4: The training and inference time (in seconds) for large CNNs and Vision Transformer architecture on the CIFAR-10 dataset.

Method	ACC1 \uparrow	ACC2 \uparrow	ACC3 \downarrow	MIS \uparrow
Gumbel-Softmax	0.950	0.946	0.910	0.240
Bernoulli-Sampling	0.951	0.948	0.333	0.243

Table 5: Performance comparison of PICNN using different strategies for generating filter-cluster assignment matrix on the CIFAR-10 dataset. ResNet-18 is used as the backbone of PICNN.

each layer in DenseNet-121 is connected to all other layers. As a result, the target conv-layer may include too many low-level concepts from early conv-layers, which leads to difficulties in grouping filters into class-specific clusters. We can also see that PICNN achieves higher ACC1 values in 14 out of 21 experiments. Besides, the performance of PICNN with ViT-b-12 as the backbone demonstrates that our approach is not only applicable to CNN architectures but also to transformer architectures.

We report the training and inference time of STD and PICNN for large CNNs and Vision Transformer architecture on CIFAR-10 dataset in Table 4. Notably, PICNN introduces a computational overhead of no more than 5% to the STD models.

Effectiveness of the Bernoulli Sampling Instead of using our Bernoulli sampling to generate the filter-cluster assignment matrix, we can apply Gumbel-Softmax to the correspondence matrix \mathbf{P} . We conducted an ablation study using Gumbel-Softmax (temperature is set to 0.01) to replace the Bernoulli sampling in PICNN. It can be seen from Table 5 that the results of the Bernoulli sampling are better than those of Gumbel-Softmax, especially using the metric ACC3. One possible explanation is that Gumbel-Softmax introduces random Gumbel noise to approximate the argmax operation for enabling gradient backpropagation. The approximation affects the results.

Effect of the filter-to-class ratio For a target conv-layer in a given CNN, we set the filter-to-class ratio to 12.8. This means that for CIFAR-10, we choose the number of filters to be 128. To exploit the effect of the filter-to-class ratio r , we vary $r = N/K$ on the CIFAR-10 dataset, where N is the number of filters in the target conv-layer and K is the number of image classes. As shown in Figure 5, PICNN shows consistent advantages over STD (ResNet-18) across

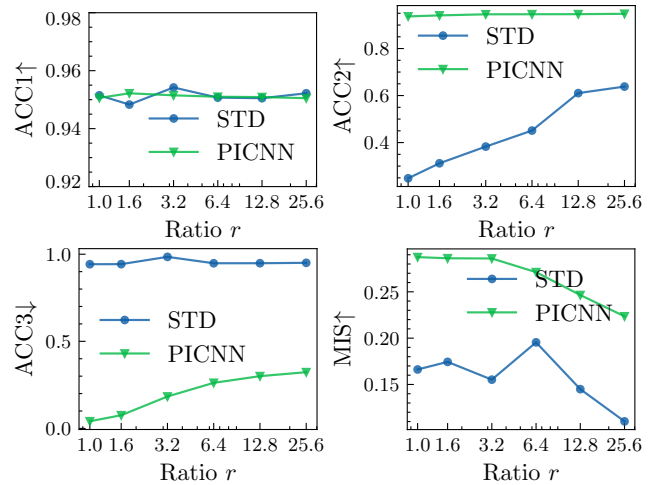


Figure 5: Effect of the filter-to-class ratio $r = N/K$ on the performance. The backbone (STD) of PICNN is ResNet-18.




Figure 6: Effect of regularization parameter λ of PICNN using ResNet-18 as the backbone on the CIFAR-10 dataset. PICNN works well across a wide range of λ .

different r values. This indicates that PICNN is robust to the filter-to-class ratio.

Effect of Regularization Parameter λ Figure 6 plots the curves of ACC1, ACC2, ACC3, and MIS of PICNN using ResNet-18 as the backbone on the CIFAR-10 dataset when varying the values of λ from 0 to 10. ACC3 decreases steadily, MIS increases fluctuately, while ACC1 and ACC2 remain stable as λ increases. This implies that with the increase of interpretation pathway weights, the filters at the target conv-layer become more class-specific and the discrimination power of the discrimination pathway is preserved.

Visualization Figure 7 displays the Grad-CAM visualizations on the PASCAL VOC Part dataset using three different sets of filters: (a) all filters, (b) a cluster of class-specific filters, and (c) the complementary clusters of class-specific fil-

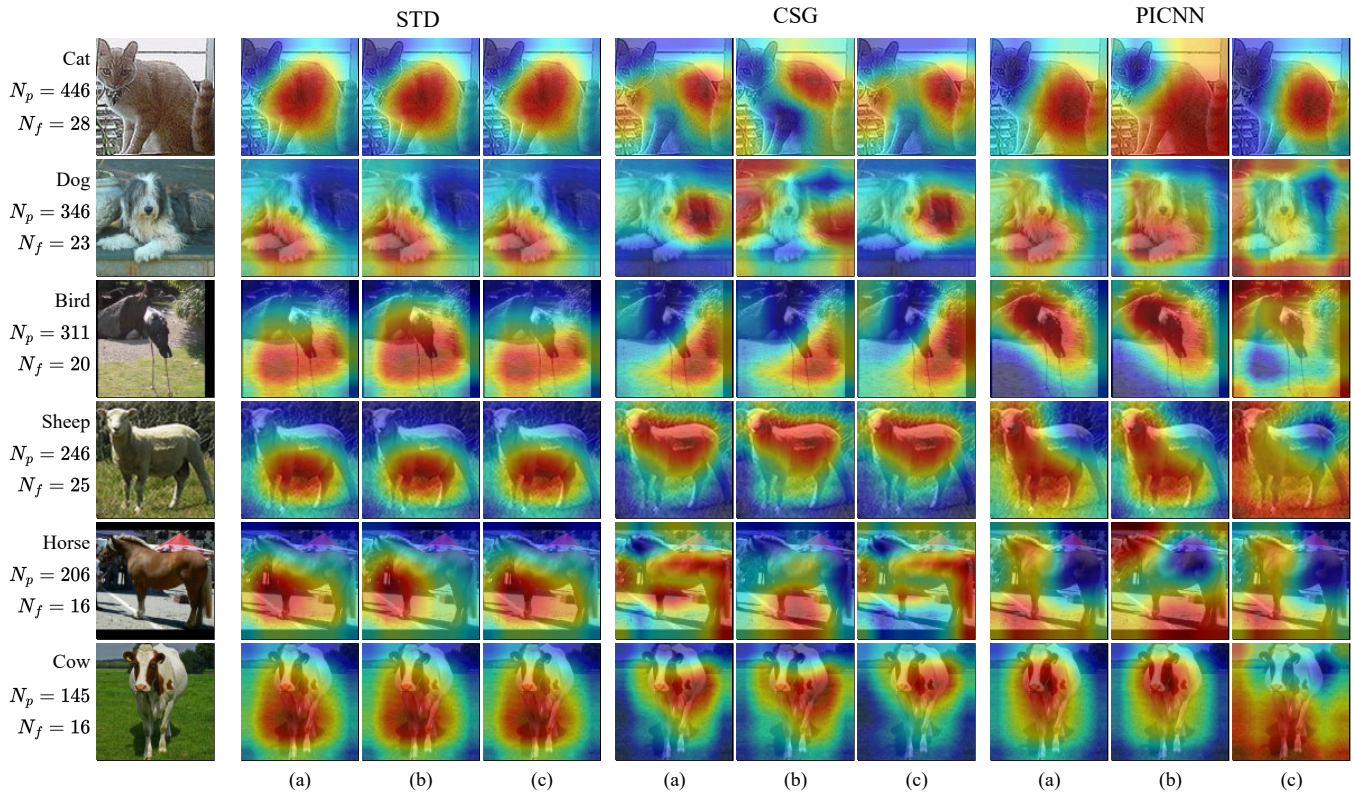


Figure 7: The Grad-CAM visualizations of filters on the PASCAL VOC Part dataset learned by STD (ResNet-18), CSG, and PICNN. CSG and PICNN use the same STD (ResNet-18) as the backbone. In each method, the (a) column represents CAMs using all filters, (b) column represents CAMs using a cluster of class-specific filters, and (c) column represents CAMs using the complementary clusters of class-specific filters. N_p is the number of pictures in a class and N_f is the number of class-specific filters assigned to a class.

ters. As above, we randomly initialize the correspondence matrix \mathbf{P} and use our Bernoulli sampling to generate the filter-cluster assignment matrix for STD (ResNet-18). We observe few differences among these three Grad-CAM visualizations in STD. This phenomenon proves that the filters and classes are entangled in STD. The Grad-CAM visualizations using the second set of filters (column (b)) of PICNN sometimes capture more class-specific information than using the first set of filters (column (a)), such as highlighting more parts of the cat’s body and the horse’s head. The Grad-CAM visualizations using the third set of filters (column (c)) of PICNN mostly highlight unimportant image regions, such as the meadow and road. We can see that the Grad-CAM visualizations using the second and third sets of filters of PICNN differ a lot. But we do not see a significant difference in STD and CSG.

Conclusion

In this paper, we have proposed a novel pathway to transform a standard CNN into an interpretable CNN without compromising its high discrimination power. The proposed pathway uses the Bernoulli sampling to generate the filter-cluster assignment matrix from a learnable filter-class

correspondence matrix. The filter-cluster assignment matrix groups the filters in the target late conv-layer in CNN into class-specific clusters and thus mitigates the filter-class entanglement problem. Because the Bernoulli sampling is non-differentiable, we propose a reparameterization trick for end-to-end learning. Experiments have shown that our method PICNN is superior to standard CNNs in terms of both interpretability and discrimination power. Moreover, our pathway has good versatility and can be combined with various network architectures.

Acknowledgments

We thank the anonymous reviewers for their valuable and constructive comments. This work was supported in part by the National Key Research and Development Program of China under Grant 2020AAA0108100.

References

- Bau, D.; Zhou, B.; Khosla, A.; Oliva, A.; and Torralba, A. 2017. Network dissection: Quantifying interpretability of deep visual representations. In *CVPR*.
- Chen, X.; Mottaghi, R.; Liu, X.; Fidler, S.; Urtasun, R.; and Yuille, A. 2014. Detect what you can: Detecting and rep-

- representing objects using holistic models and body parts. In *CVPR*.
- Coates, A.; Ng, A.; and Lee, H. 2011. An analysis of single-layer networks in unsupervised feature learning. In *AISTATS*.
- D’Amour, A.; Heller, K.; Moldovan, D.; Adlam, B.; Alipanahi, B.; Beutel, A.; Chen, C.; Deaton, J.; Eisenstein, J.; Hoffman, M. D.; et al. 2022. Underspecification presents challenges for credibility in modern machine learning. *JMLR*.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *CVPR*.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale.
- Dosovitskiy, A.; and Brox, T. 2016. Inverting visual representations with convolutional networks. In *CVPR*.
- Fong, R.; and Vedaldi, A. 2018. Net2vec: Quantifying and explaining how concepts are encoded by filters in deep neural networks. In *CVPR*.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. In *NeurIPS*.
- Hasany, S. N.; Petitjean, C.; and Mériaudeau, F. 2023. Seg-XXRes-CAM: Explaining Spatially Local Regions in Image Segmentation. In *CVPR*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *CVPR*.
- Hendricks, L. A.; Akata, Z.; Rohrbach, M.; Donahue, J.; Schiele, B.; and Darrell, T. 2016. Generating visual explanations. In *ECCV*.
- Huang, G.; Liu, Z.; Van Der Maaten, L.; and Weinberger, K. Q. 2017. Densely connected convolutional networks. In *CVPR*.
- Jang, E.; Gu, S.; and Poole, B. 2016. Categorical Reparameterization with Gumbel-Softmax. In *ICLR*.
- Jiang, Z.; Wang, Y.; Davis, L.; Andrews, W.; and Rozgic, V. 2017. Learning discriminative features via label consistent neural network. In *WACV*.
- Kingma, D. P.; and Welling, M. 2014. Auto-Encoding Variational Bayes. In *ICLR*.
- Krizhevsky, A. 2014. One weird trick for parallelizing convolutional neural networks. *CoRR*.
- Krizhevsky, A.; Hinton, G.; et al. 2009. Learning multiple layers of features from tiny images.
- Kweon, H.; Yoon, S.-H.; Kim, H.; Park, D.; and Yoon, K.-J. 2021. Unlocking the potential of ordinary classifier: Class-specific adversarial erasing framework for weakly supervised semantic segmentation. In *ICCV*.
- Lee, J. R.; Kim, S.; Park, I.; Eo, T.; and Hwang, D. 2021. Relevance-cam: Your model already knows where to look. In *CVPR*.
- Li, Y.; Li, Y.; Lu, J.; Shechtman, E.; Lee, Y. J.; and Singh, K. K. 2021. Collaging class-specific gans for semantic image synthesis. In *ICCV*.
- Liang, H.; Ouyang, Z.; Zeng, Y.; Su, H.; He, Z.; Xia, S.-T.; Zhu, J.; and Zhang, B. 2020. Training interpretable convolutional neural networks by differentiating class-specific filters. In *ECCV*.
- Mahendran, A.; and Vedaldi, A. 2015. Understanding deep image representations by inverting them. In *CVPR*.
- Martinez, B.; Modolo, D.; Xiong, Y.; and Tighe, J. 2019. Action recognition with spatial-temporal discriminative filter banks. In *ICCV*.
- Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; et al. 2019. Pytorch: An imperative style, high-performance deep learning library. In *NeurIPS*.
- Rudin, C. 2019. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nat. Mac. Intell.*
- Sandler, M.; Howard, A.; Zhu, M.; Zhmoginov, A.; and Chen, L.-C. 2018. Mobilenetv2: Inverted residuals and linear bottlenecks. In *CVPR*.
- Sarkar, S.; Babu, A. R.; Gundecha, V.; Guillen, A.; Mousavi, S.; Luna, R.; Ghorbanpour, S.; and Naug, A. 2023. RL-CAM: Visual Explanations for Convolutional Networks Using Reinforcement Learning. In *CVPR*.
- Selvaraju, R. R.; Cogswell, M.; Das, A.; Vedantam, R.; Parikh, D.; and Batra, D. 2017. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *ICCV*.
- Shen, W.; Wei, Z.; Huang, S.; Zhang, B.; Fan, J.; Zhao, P.; and Zhang, Q. 2021. Interpretable compositional convolutional neural networks. In *IJCAI*.
- Shi, J.; and Malik, J. 2000. Normalized cuts and image segmentation. *TPAMI*.
- Simonyan, K.; Vedaldi, A.; and Zisserman, A. 2013. Deep inside convolutional networks: Visualising image classification models and saliency maps. In *ICLR*.
- Simonyan, K.; and Zisserman, A. 2014. Very deep convolutional networks for large-scale image recognition. In *ICLR*.
- Springenberg, J. T.; Dosovitskiy, A.; Brox, T.; and Riedmiller, M. 2014. Striving for simplicity: The all convolutional net. In *ICLR*.
- Tan, M.; and Le, Q. 2019. Efficientnet: Rethinking model scaling for convolutional neural networks. In *ICML*.
- Tang, H.; Xu, D.; Yan, Y.; Torr, P. H.; and Sebe, N. 2020. Local class-specific and global image-level generative adversarial networks for semantic-guided scene generation. In *CVPR*.
- Wang, J.; Liu, H.; Wang, X.; and Jing, L. 2021. Interpretable image recognition by constructing transparent embedding space. In *ICCV*.
- Wang, Y.; Morariu, V. I.; and Davis, L. S. 2018. Learning a discriminative filter bank within a cnn for fine-grained recognition. In *CVPR*.

- Yang, Y.; Kim, S.; and Joo, J. 2022. Explaining deep convolutional neural networks via latent visual-semantic filter attention. In *CVPR*.
- Yosinski, J.; Clune, J.; Nguyen, A.; Fuchs, T.; and Lipson, H. 2015. Understanding neural networks through deep visualization.
- Zablocki, É.; Ben-Younes, H.; Pérez, P.; and Cord, M. 2022. Explainability of deep vision-based autonomous driving systems: Review and challenges. *IJCV*.
- Zeiler, M. D.; and Fergus, R. 2014. Visualizing and understanding convolutional networks. In *ECCV*.
- Zerhouni, E.; Lányi, D.; Viana, M.; and Gabrani, M. 2017. Wide residual networks for mitosis detection. In *ISBI*. IEEE.
- Zhang, Q.; Cao, R.; Shi, F.; Wu, Y. N.; and Zhu, S.-C. 2018. Interpreting CNN knowledge via an explanatory graph. In *AAAI*.
- Zhang, Q.; Wu, Y. N.; and Zhu, S.-C. 2018. Interpretable convolutional neural networks. In *CVPR*.
- Zhang, Q.; Yang, Y.; Ma, H.; and Wu, Y. N. 2019. Interpreting cnns via decision trees. In *CVPR*.
- Zhou, B.; Khosla, A.; Lapedriza, À.; Oliva, A.; and Torralba, A. 2015. Object Detectors Emerge in Deep Scene CNNs. In *ICLR*.
- Zhou, B.; Khosla, A.; Lapedriza, A.; Oliva, A.; and Torralba, A. 2016. Learning deep features for discriminative localization. In *CVPR*.