

# Learning Bottleneck Concepts in Image Classification

Bowen Wang<sup>1</sup>, Liangzhi Li<sup>2\*</sup>, Yuta Nakashima<sup>2</sup>, Hajime Nagahara<sup>2</sup>  
 Osaka University, Japan

<sup>1</sup>bowen.wang@is.ids.osaka-u.ac.jp

<sup>2</sup>{li, n-yuta, nagahara}@ids.osaka-u.ac.jp

## Abstract

Interpreting and explaining the behavior of deep neural networks is critical for many tasks. Explainable AI provides a way to address this challenge, mostly by providing per-pixel relevance to the decision. Yet, interpreting such explanations may require expert knowledge. Some recent attempts toward interpretability adopt a concept-based framework, giving a higher-level relationship between some concepts and model decisions. This paper proposes *Bottleneck Concept Learner (BotCL)*, which represents an image solely by the presence/absence of concepts learned through training over the target task without explicit supervision over the concepts. It uses self-supervision and tailored regularizers so that learned concepts can be human-understandable. Using some image classification tasks as our testbed, we demonstrate *BotCL*'s potential to rebuild neural networks for better interpretability<sup>1</sup>.

## 1. Introduction

Understanding the behavior of deep neural networks (DNNs) is a major challenge in the explainable AI (XAI) community, especially for medical applications [19, 38], for identifying biases in DNNs [2, 18, 42], *etc.* Tremendous research efforts have been devoted to the post-hoc paradigm for a *posteriori* explanation [29, 33]. This paradigm produces a relevance map to spot regions in the input image that interact with the model's decision. Yet the relevance map only tells *low-level* (or per-pixel) relationships and does not explicitly convey any semantics behind the decision. Interpretation of relevance maps may require expert knowledge.

The *concept-based* framework [22, 37, 50] is inspired by the human capacity to learn a new concept by (subconsciously) finding finer-grained concepts and reuse them in different ways for better recognition [24]. Instead of giving per-pixel relevance, this framework offers higher-level

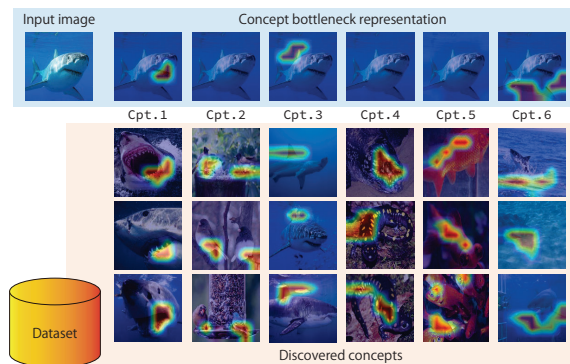


Figure 1. Examples of concepts discovered by BotCL in ImageNet [10] and concepts in the input image. BotCL automatically discovers a set of concepts optimized for the target task and represents an image solely with the presence/absence of concepts.

relationships between the image and decision mediated by a limited number of *concepts*. That is, the decision is explained by giving a set of concepts found in the image. The interpretation of the decision is thus straightforward once the interpretation of each concept is established.

Some works use concepts for the post-hoc paradigm for better interpretation of the decision [14, 50], while the link between the decision and concepts in the image is not obvious. The *concept bottleneck* structure [23] uses the presence/absence of concepts as image representation (referred to as *concept activation*). The classifier has access only to the concept activation, so the decision is strongly tied to the concepts. This bottleneck structure has become the mainstream of the concept-based framework [5, 20, 28, 31].

A major difficulty in this framework is designing a set of concepts that suits the target task. A promising approach is handcrafting them [4, 21, 48], which inherently offers better interpretability at the cost of extra annotations on the concepts. Recent attempts automatically discover concepts [1, 13, 14, 46]. Such concepts may not always be consistent with how humans (or models) see the world [25, 47] and may require some effort to interpret them, but concept discovery without supervision is a significant advantage.

\*Corresponding author.

<sup>1</sup>Code is available at <https://github.com/wbw520/BotCL> and a simple demo is available at <https://botcl.liangzhili.com/>.

Inspired by these works, we propose **bottleneck concept learner** (BotCL) for simultaneously discovering concepts and learning the classifier. BotCL optimizes concepts for the given target image classification task without supervision for the concepts. An image is represented solely by the existence of concepts and is classified using them. We adopt a slot attention-based mechanism [26, 27] to spot the region in which each concept is found. This gives an extra signal for interpreting the decision since one can easily see what each learned concept represents by collectively showing training images with the detected concepts. Figure 1 shows examples from ImageNet [10]. BotCL discovers a predefined number of concepts in the dataset, which are exemplified by several images with attention maps. An image of *Great White Shark* is represented by the right part of mouth (Cpt. 1) and fins (Cpt. 3). BotCL uses a single fully-connected (FC) layer as a classifier, which is simple but enough to encode the co-occurrence of each concept and each class.

**Contribution.** For better concept discovery, we propose to use *self-supervision over concepts*, inspired by the recent success in representation learning [9, 16]. Our ablation study demonstrates that self-supervision by contrastive loss is the key. We also try several constraints on concepts themselves, *i.e.*, *individual consistency* to make a concept more selective and *mutual distinctiveness* for better coverage of various visual elements. These additional constraints regular the training process and help the model learn concepts of higher quality.

## 2. Related Works

### 2.1. Explainable AI

XAI focuses on uncovering black-box deep neural networks [3, 6, 12, 32, 35, 36, 41, 43, 46]. A major approach is generating a relevance map that spots important regions for the model’s decision. Various methods have been designed for specific architectures, *e.g.*, CAM [49], and Grad-CAM [33] for convolutional neural networks; [7] for Transformers [40]. However, the interpretation of the relevance maps may not always be obvious, which spurs different approaches [34, 45], including context-based ones.

### 2.2. Concept-based framework for interpretability

A straightforward way to define a set of concepts for a target task is to utilize human knowledge [22, 48]. Such concepts allow quantifying their importance for a decision [21]. A large corpus of concepts [4, 39] is beneficial for delving into hidden semantics in DNNs [50]. These methods are of the post-hoc XAI paradigm, but a handcrafted set of concepts can also be used as additional supervision for models with the concept bottleneck structure [15, 22, 31].

Handcrafting a set of concepts offers better interpretability

as they suit human perception; however, the annotation cost is non-negligible. Moreover, such handcrafted concepts may not always be useful for DNNs [47]. These problems have motivated automatic concept discovery. Superpixels are a handy unit for finding low-level semantics, and concepts are defined by clustering them [13, 14, 30]. Another interesting approach is designing a set of concepts to be sufficient statistics of original DNN features [46]. These methods are designed purely for interpretation, and concept discovery is made aside from training on the target task.

The concept bottleneck structure allows optimizing a set of concepts for the target task. ProtoPNet [8] adopts this structure and identifies concepts based on the distance between features and concepts. SENN [1] uses self-supervision by reconstruction loss for concept discovery.

SENN inspired us to use self-supervision, but instead of reconstruction loss, we adopt contrastive loss tailored. For a natural image classification task, this contrastive loss is essential for concept discovery.

## 3. Model

Given a dataset  $\mathcal{D} = \{(x_i, y_i) | i = 1, 2, \dots, N\}$ , where  $x_i$  is an image and  $y_i$  is the target class label in the set  $\Omega$  associated with  $x_i$ . BotCL learns a set of  $k$  concepts while learning the original classification task. Figure 2a shows an overview of BotCL’s training scheme, consisting of a concept extractor, regularizers, and a classifier, as well as self-supervision (contrastive and reconstruction losses).

For a new image  $x$ , we extract feature map  $F = \Phi(x) \in \mathbb{R}^{d \times h \times w}$  using a backbone convolutional neural network  $\Phi$ .  $F$  is then fed into the concept extractor  $g_C$ , where  $C$  is a matrix, each of whose  $\kappa$ -th column vector  $c_\kappa$  is a *concept prototype* to be learned. The concept extractor produces concept bottleneck activations  $t \in [0, 1]^k$ , indicating the presence of each concept, as well as concept features  $V \in \mathbb{R}^{d \times k}$  from regions where each concept exists. The concept activations in  $t$  are used as input to the classifier to compute score  $s \in [0, 1]^{|\Omega|}$ . We use self-supervision and regularizers for training, taking  $t$  and  $V$  as input to constrain the concept prototypes.

### 3.1. Concept Extractor

Concept extractor uses slot attention [26, 27]-based mechanism to discover visual concepts in  $\mathcal{D}$ . We first add position embedding  $P$  to feature map  $F$  to retain the spatial information, *i.e.*,  $F' = F + P$ . The spatial dimension of  $F'$  is flattened, so its shape is  $l \times d$ , where  $l = hw$ .

The slot-attention computes attention over the spatial dimension for concept  $\kappa$  from  $c_\kappa$  and  $F'$ . Let  $Q(c_\kappa) \in \mathbb{R}^d$ , and  $K(F') \in \mathbb{R}^{d \times l}$  denote nonlinear transformations for  $c_\kappa$  and  $F'$ , respectively, given as multi-layer perceptrons with three FC layers and a ReLU nonlinearity between them. Attention  $a_\kappa \in [0, 1]^l$  is given using a normalization function




Figure 2. (a) The model pipeline. (b) Self-supervision and regularizers.

$\phi$  (refer to **supp. material**) as

$$a_\kappa = \phi(Q(c_\kappa)^\top K(F')). \quad (1)$$

This attention indicates where concept  $\kappa$  presents in the image as shown in Figure 1. If concept  $\kappa$  is absent, corresponding entries of  $a_\kappa$  are all close to 0. We summarize the presence of each concept into concept activation  $t_\kappa$  by reducing the spatial dimension of  $a_\kappa$  as  $t_\kappa = \tanh(\sum_m a_{\kappa m})$ , where  $a_{\kappa m}$  is the  $m$ -th element of  $a_\kappa$ .

### 3.2. Feature Aggregation

For training, we also aggregate features in  $F$  corresponding to concept  $\kappa$  into concept feature  $v_\kappa$  by

$$v_\kappa = F a_\kappa, \quad (2)$$

which gives the average of image features over the spatial dimension weighted by attention.

### 3.3. Classifier

We use a single FC layer without a bias term as the classifier, and concept activation  $t = (t_1, \dots, t_k)^\top$  is the only input, serving as the concept bottleneck [22]. Formally, letting  $W$  be a learnable matrix, prediction  $\hat{y} \in \mathbb{R}^{|\Omega|}$  is given by

$$\hat{y} = Wt. \quad (3)$$

This classifier can be roughly interpreted as learning the correlation between the class and concepts. Let  $w_\omega$  be the raw vector of  $W$  corresponding to class  $\omega \in \Omega$ , and  $w_{\omega\kappa}$  is its  $\kappa$ -th element. A positive value of  $w_{\omega\kappa}$  means that concept  $\kappa$  co-occurs with class  $\omega$  in the dataset, so its presence in a new image positively supports class  $\omega$ . Meanwhile, a negative value means the concept rarely co-occurs.

## 4. Training

### 4.1. Self-supervision for Concept Discovery

The absence of concept labels motivates us to incorporate self-supervision for concept discovery. We employ two losses for different types of target tasks.

**Reconstruction loss.** SENN [1] uses an autoencoder-like structure for learning better representation. We assume this structure works well when visual elements are strongly tied with the position<sup>2</sup> since even discrete concepts should have sufficient information to reconstruct the original image. Based on this assumption, we design a reconstruction loss for self-supervision. As shown in Figure 2b, decoder  $D$  only takes  $t$  as input and reconstructs the original image. We define our reconstruction loss as

$$l_{\text{rec}} = \frac{1}{|\mathcal{B}|} \sum_{x \in \mathcal{B}} \|D(t) - x\|^2. \quad (4)$$

**Contrastive loss.** The composition of natural images is rather arbitrary, so information in  $t$  should be insufficient to reconstruct the original image. we thus design a simple loss for an alternative, borrowing the idea from the recent success of contrastive learning for self-supervision [9, 16].

We leverage the image-level labels of the target classification task. Let  $\hat{t} = 2t - \mathbf{1}_k$ , where  $\mathbf{1}_k$  is the  $k$ -dimensional vector with all elements being 1. If a pair  $(\hat{t}, \hat{t}')$  of concept activations belong to the same class (*i.e.*,  $y = y'$  for  $y$  and  $y'$  corresponding to  $\hat{t}$  and  $\hat{t}'$ ), they should be similar to each other since a similar set of concepts should be in the corresponding images, and otherwise dissimilar. The number  $|\Omega|$  of classes can be smaller than the number  $|\mathcal{B}|$  of images in a mini-batch so that a mini-batch can have multiple images of the same class. Therefore, we use sigmoid instead of softmax, leading to

$$l_{\text{ret}} = -\frac{1}{|\mathcal{B}|} \sum \alpha(y, y') \log J(\hat{t}, \hat{t}', y, y'), \quad (5)$$

where  $\alpha$  is the weight to mitigate the class imbalance problem (see **supp. material**) and

$$J(\hat{t}, \hat{t}', y, y') = \begin{cases} \sigma(\hat{t}^\top \hat{t}') & \text{for } y = y' \\ 1 - \sigma(\hat{t}^\top \hat{t}') & \text{otherwise} \end{cases}. \quad (6)$$

<sup>2</sup>For example, images of “7” in MNIST almost always have the acute angle in the top-right part.

## 4.2. Concept Regularizers

We also employ concept regularizers to facilitate training. They constrain concept prototypes  $\{c_\kappa\}$  through  $\{v_\kappa\}$ .

**Individual consistency.** For better interpretability, each learned concept should not have large variations. That is, the concept features  $v_\kappa$  and  $v'_\kappa$  of different images should be similar to each other if  $t_\kappa$  is close to 1. Let  $\mathcal{H}_\kappa$  denote the set of all concept features of different images in a mini-batch, whose activation is larger than the empirical threshold  $\xi$ , which is dynamically calculated as the mean of  $t_\kappa$  in a mini-batch. Using the cosine similarity  $\text{sim}(\cdot, \cdot)$ , we define the consistency loss as:

$$l_{\text{con}} = -\frac{1}{k} \sum_{\kappa} \sum_{v_\kappa, v'_\kappa} \frac{\text{sim}(v_\kappa, v'_\kappa)}{|\mathcal{H}_\kappa|(|\mathcal{H}_\kappa| - 1)}, \quad (7)$$

where the second summation is computed over all combinations of concept features  $v_\kappa$  and  $v'_\kappa$ . This loss penalizes a smaller similarity between  $v_\kappa$  and  $v'_\kappa$ .

**Mutual distinctiveness.** To capture different aspects of images, different concepts should cover different visual elements. This means that the average image features of concept  $\kappa$  within a mini-batch, given by  $\bar{v}_\kappa = \sum_{v_\kappa \in \mathcal{H}_\kappa} v_\kappa$ , should be different from any other  $v_{\kappa'}$ . We can encode this into a loss term as

$$l_{\text{dis}} = \sum_{\kappa, \kappa'} \frac{\text{sim}(\bar{v}_\kappa, \bar{v}_{\kappa'})}{k(k-1)}, \quad (8)$$

where the summation is computed over all combinations of concepts. Note that concept  $\kappa$  is excluded from this loss if no image in a mini-batch has concept  $\kappa$ .

## 4.3. Quantization Loss

Concept activation  $t$  can be sufficiently represented by a binary value, but we instead use a continuous value for training. We design a quantization loss to guarantee values are close to 0 or 1, given by

$$l_{\text{qua}} = \frac{1}{k|\mathcal{B}|} \sum_{x \in \mathcal{B}} \|\text{abs}(\hat{t}) - \mathbf{1}_\kappa\|^2, \quad (9)$$

where  $\text{abs}(\cdot)$  gives the element-wise absolute value and  $\|\cdot\|$  gives the Euclidean norm.

## 4.4. Total Loss

We use softmax cross-entropy for the target classification task's loss, denoted by  $l_{\text{cls}}$ . The overall loss of BotCL is defined by combining the losses above as

$$L = l_{\text{cls}} + \lambda_{\text{R}} l_{\text{R}} + \lambda_{\text{con}} l_{\text{con}} + \lambda_{\text{dis}} l_{\text{dis}} + \lambda_{\text{qua}} l_{\text{qua}}, \quad (10)$$

where  $l_{\text{R}}$  is either  $l_{\text{rec}}$  or  $l_{\text{ret}}$  depending on the target domain,  $\lambda_{\text{qua}}$ ,  $\lambda_{\text{con}}$ ,  $\lambda_{\text{dis}}$ , and  $\lambda_{\text{R}}$  are weights to balance each term.

## 5. Results

### 5.1. Experimental Settings

We evaluate BotCL on MNIST [11], CUB200 [44], and ImageNet [10]. For evaluating discovered concepts, we re-generated a synthetic shape dataset (Synthetic) [46].

For MNIST, we applied the same networks as [1] for the backbone and the concept decoder. For CUB200 (same data split as [22]) and ImageNet, we used pre-trained ResNet [17] as the backbone with a  $1 \times 1$  convolutional layer to reduce the channel number (512 for ResNet-18 and 2048 for ResNet-101) to 128. We chose a concept number  $k = 20$  for MNIST and  $k = 50$  for the other natural image datasets. To generate Synthetic, we followed the setting of [46], where 18,000 images were generated for training and 2,000 for evaluation. We used  $k = 15$  with ResNet-18 backbone.

Images were resized to  $256 \times 256$  and cropped to  $224 \times 224$  (images in Synthetic were directly resized to  $224 \times 224$ ). Only random horizontal flip was applied as data augmentation during training. The weight of each loss was defaulted to  $\lambda_{\text{qua}} = 0.1$ ,  $\lambda_{\text{con}} = 0.01$ ,  $\lambda_{\text{dis}} = 0.05$ , and  $\lambda_{\text{R}} = 0.1$ .

### 5.2. Classification Performance

We compare the performance of BotCL with corresponding baselines (LeNet for MNIST and ResNet-18 for others with a linear classifier), our reimplement of k-means and PCA in [46],<sup>3</sup> and state-of-the-art concept-based models. Table 1 summarized the results. BotCL with contrastive loss (BotCL<sub>Cont</sub>) achieves the best accuracy on CUB200, ImageNet, and Synthetic, outperforming the baseline linear classifiers. It is also comparable to the state-of-the-art on MNIST and Synthetic. BotCL with reconstruction loss (BotCL<sub>Rec</sub>) shows a performance drop over CUB200, ImageNet, and Synthetic, while it outperforms BotCL<sub>Cont</sub> on MNIST. This behavior supports our assumption that the reconstruction loss is useful only when concepts are strongly tied to their spatial position. Otherwise,  $t$  is insufficient to reconstruct the original image, and BotCL fails. Contrastive self-supervision is the key to facilitating concept discovery.

We also explore the relationship between the number of classes and BotCL's accuracy over CUB200 and ImageNet. We used small and large variants of ResNet as the backbone. We extracted subsets of the datasets consisting of the first  $n$  classes along with the class IDs. Figure 3 shows that BotCL has a competitive performance when the number of classes is less than 200. We conclude that BotCL hardly degrades the classification performance on small- or middle-sized datasets. However, this is not the case for  $n > 200$  (refer to **supp. material** for larger  $n$  and different  $k$ 's).

<sup>3</sup>Implementation details are in **supp. material**.






Figure 3. Classification accuracy vs. the number of classes. We used subsets of CUB200 and ImageNet with  $k = 50$  and ResNet-18 and ResNet-101 backbones.

Table 1. Performance comparison in classification accuracy. The best concept-based method is highlighted in bold. BotCL<sub>Rec</sub> and BotCL<sub>Cont</sub> are both BotCL but with reconstruction and contrastive loss, respectively. For ImageNet, we used the first 200 classes.

	CUB200	ImageNet	MNIST	Synthetic
Baseline	0.731	0.786	0.988	0.999
k-means* [46]	0.063	0.427	0.781	0.747
PCA* [46]	0.044	0.139	0.653	0.645
SENN [1]	0.642	0.673	<b>0.985</b>	0.984
ProtoPNet [8]	0.725	0.752	0.981	0.992
BotCL <sub>Rec</sub>	0.693	0.720	0.983	0.785
BotCL <sub>Cont</sub>	<b>0.740</b>	<b>0.795</b>	0.980	<b>0.998</b>

### 5.3. Interpretability

#### 5.3.1 Qualitative validation of discovered concepts

Figure 4a visualizes  $a_{\kappa}$ , showing concept  $\kappa$  in the image, over MNIST. We selected 5 concepts out of 20 that are most frequently activated (*i.e.*,  $t_{\kappa} > 0.5$ ) in the training set.<sup>4</sup> Taking digits 0 and 9 as an example, we can observe that they share Cpts. 3–5 and the only difference is Cpt. 2 that locates in the lower edge of the vertical stroke of 9. This stroke is specific to digit 9. We used BotCL<sub>Rec</sub>, so we can remove Cpt. 2 before reconstruction, which generates an image like 0 (refer to Section 5.3.3). Some concepts are incompatible with human intuition; yet we can interpret such concepts (*e.g.*, Cpt. 1 may attend to the missing stroke that completes a circle).

For CUB200, we train BotCL<sub>Const</sub> with  $n = 50$  and  $k = 20$ . Figure 5a shows the attention maps of an image of yellow headed black bird. We can observe that the attentions for Cpts. 1–5 cover different body parts, including the head, wing, back, and feet, which proves that BotCL can learn valid concepts from the natural image as well. **Supp. material** exemplifies all concepts discovered from MNIST and CUB200.

<sup>4</sup>Cpts. 1–5 are ordered based on the frequency counted in the dataset.

#### 5.3.2 Consistency and distinctiveness of each concept

BotCL is designed to discover individually consistent and mutually distinctive concepts. We qualitatively verify this by showing each concept with its top-5 activated images<sup>5</sup> with attention maps in Figure 4b. For MNIST, different concepts cover different patterns, and each concept covers the same patterns in different samples (even the samples of different classes). Figure 5b for CUB200 shows that BotCL renders a similar behavior on the CUB200 dataset; that is, the top-5 concepts are responsible for different patterns, and each of them is consistent.

#### 5.3.3 Contribution of each concept in inference

We can qualitatively see the contribution of each concept by removing the concept and seeing the changes in the corresponding self-supervision task’s output. As shown in Figure 4c, when we set the activation of Cpt. 2 (responsible for the vertical stroke of digit 9) to zero, the reconstructed image looks like digit 0. When Cpt. 1, representing the absence of the circle in digit 7, is deactivated (*i.e.*,  $t_1$  is set to 0), a circle emerges in the upper part of the reconstructed image. The resulting image looks more like digit 9.

For CUB200 shown in Figure 5c, we show images most similar to the input image in Figure 5a among the dataset in terms of  $\hat{t}^T \hat{t}'$ , with ablating each concept. When Cpt. 1 (responsible for the yellow head) is deactivated, more black-head bird images appear in the top-8 images. Cpt. 5 covers birds’ feet and is common among most bird classes. Deactivating this concept does not change the top-8 images. These observations suggest that although some concepts do not contribute to classification performance, images are successfully represented by combinations of concepts.

### 5.4. Quantitative Evaluation on Synthetic

One problem of the concept-based approach is the absence of established quantitative evaluations of concepts because the choice of concepts may be arbitrary and the same level of representability may be achieved with different sets of concepts. A single predefined set of concepts

<sup>5</sup>For each concept  $\kappa$ , five images whose  $t_{\kappa}$  is highest among  $\mathcal{D}$ .

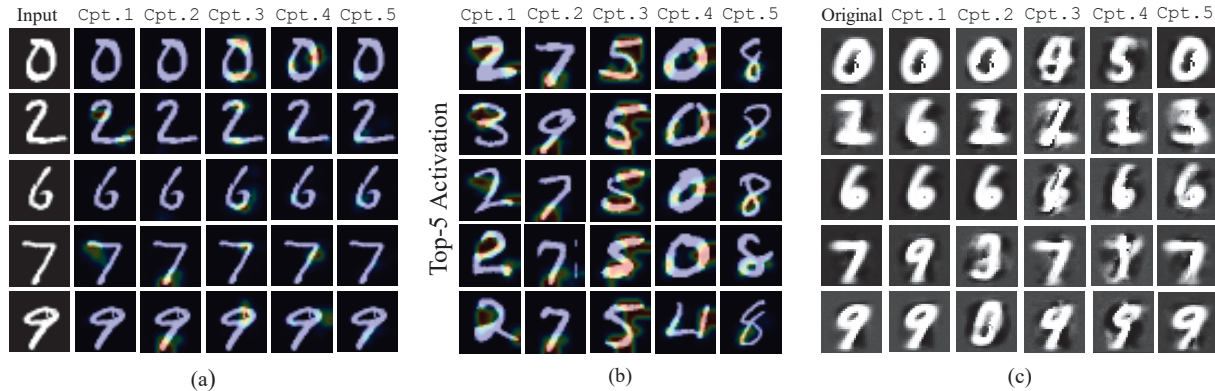


Figure 4. Concepts for MNIST. (a) Attention maps for different input images. (b) Top-5 activated images (images in the dataset whose  $t_\kappa$  is largest) for each concept. (c) Images reconstructed by our concept decoder with all detected concepts (original) and with a certain concept deactivated.

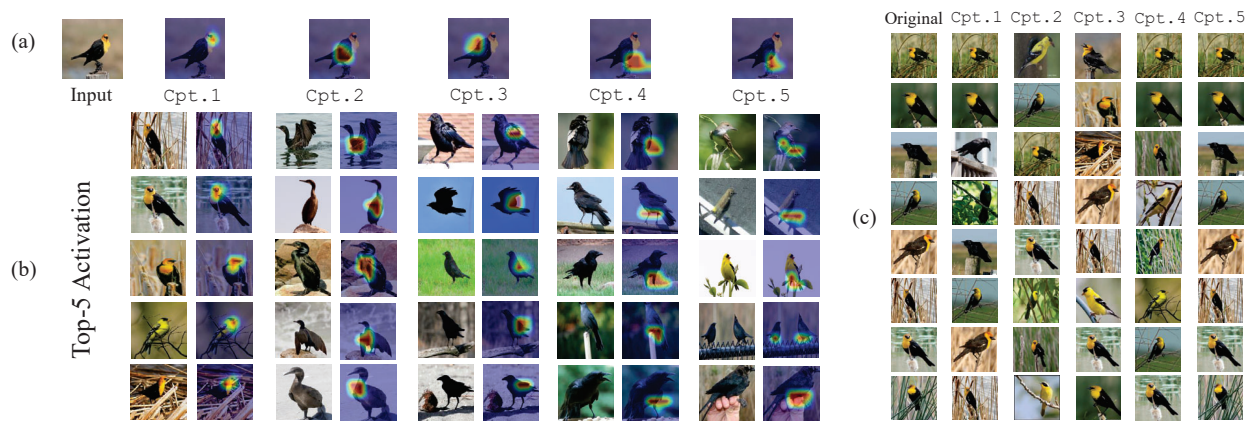


Figure 5. Concepts learned for CUB200. (a) Visualization of 5 most important concepts for yellow headed black bird. (b) Top-5 activated concepts. (c) Image retrieval when all detected concepts were used (original) and when a certain concept was deactivated.

is not enough to evaluate the goodness of discovered concepts. Literature has evaluated concepts qualitatively (as Section 5.3) or by user study (as Section 5.5).

We decided to use Synthetic [46] for quantitatively evaluating concepts. The task is a multi-label classification that involves 15 shapes. Combinations of the 5 shapes (shown in Figure 6a, S. 1 to S. 5) form 15 classes, and the other 10 shapes are noises<sup>6</sup>. We deem a shape is covered by concept  $\kappa$  when the shape’s area and concept  $\kappa$ ’s area (the area with  $a_\kappa > \gamma$  for BotCL, where  $\gamma = 0.9$  is a predefined threshold) overlap. Let  $h_{s\kappa} = 1$  denote shape  $s$  overlaps with concept  $\kappa$ , and  $h_{s\kappa} = 0$  otherwise. The coverage of  $s$  by concept  $\kappa$  is given by

$$\text{Coverage}_{s\kappa} = \mathbb{E}[h_{s\kappa}], \quad (11)$$

where the expectation is computed over the images in the

<sup>6</sup>Images are generated with random shapes, so there can be multiple classes (combinations of shapes) in a single image, which forms a multi-label classification task.

test set with concept  $\kappa$  activated. The concepts and the 5 shapes are associated as a combinatorial optimization problem so that the sum of  $\text{Coverage}_{s\kappa}$  over  $s$  are maximized.

We use  $k = 15$  to train BotCL. Figure 6a visualizes the concept associated with each shape<sup>7</sup>. A concept is located by  $a_\kappa$  for 6 images with the highest concept activations  $t_\kappa$ . The concepts cover the associated shapes with relatively small regions, but one concept usually covers multiple shapes. This can be further evident in Figure 6b that shows  $\text{Coverage}_{s\kappa}$ . Cpt. 8 only covers S. 3, whereas Cpt. 1 and Cpt. 13 covers multiple shapes.

We use three metrics other than accuracy to evaluate the performance of concept discovery<sup>8</sup>: (i) **Completeness** measures how well a concept covers its associated shape in the dataset. (ii) **Purity** shows the ability to discover concepts that only cover a single shape. (iii) **Distinctiveness** quanti-

<sup>7</sup>Note that in this experiment, only shapes matter but not colors.

<sup>8</sup>Further details are in **supp. material**.

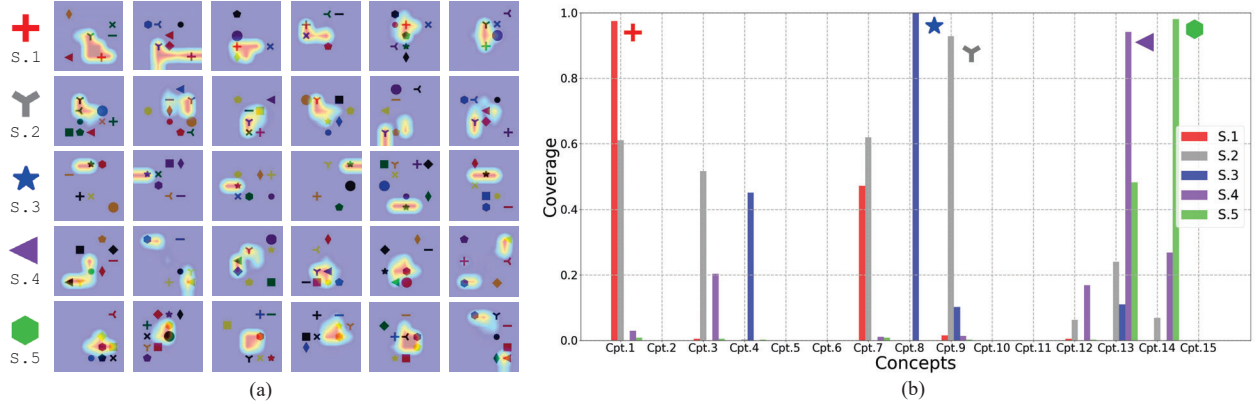


Figure 6. Experiment on Synthetic. (a) S.1-S.5 are the five shapes of which combinations form classes. Attention maps next to each shape are of the concept that covers the shape. (b) Coverage<sub>s $\kappa$</sub>  (the concept associated with each of the five shapes is marked).

fies the difference among concepts based on the coverage.

BotCL with contrastive loss is compared<sup>9</sup> with ACE [14], and two baselines PCA and k-means in [46]. We apply k-means or PCA to  $F$  of all images in the dataset after flattening the spatial dimensions. The cluster centers or the principal components are deemed as concepts. Attention maps can be computed by Euclidean distance or cosine similarity. Once the attention maps are obtained, we follow BotCL’s process for classification.

As shown in Table 2, BotCL shows better completeness, distinctiveness, and accuracy scores than comparative methods. Although k-means is able to discover concepts, they are not optimized for the target classification task, and the performance is low. As we discussed, the concepts learned by BotCL tend to cover more than one target shape, causing a comparatively low purity. The cluster center of k-means is able to capture only one kind of shape at the cost of completeness. We can also observe that all methods are affected by concept number  $k$ , and generally a larger  $k$  ensure better performance on all metrics. This result is not surprising, but we confirmed that a larger  $k$  is preferable for better interpretability. We detail the generation of the dataset, implementation of PCA and k-means, and formal definitions of metrics in **supp. material**.

### 5.5. User Study

Our user study is designed to evaluate BotCL with realistic datasets for the challenge of human understanding. Participants are asked to observe the test images with the attention map for concept  $\kappa$  (refer to Section 5.3.2) and select some phrases in the predefined vocabulary that best describes the concept (*i.e.*, attended regions). They can also choose *None of them* if they cannot find any consistent visual elements. We recruited 20 participants for each concept

<sup>9</sup>SENN [1] and ProtoNet [8] are not comparable. SENN’s concepts globally cover a whole image. ProtoNet requires way more concepts.

Table 2. Quantitative evaluation on Synthetic. Note that ACE uses concepts for post-hoc explanation and does not use them for classification. Comp., Dist., and Acc. mean completeness, distinctiveness, and accuracy, respectively.

		Comp.	Purity	Dist.	Acc.
$k = 5$	ACE	0.662	0.274	0.084	—
	k-means	0.630	0.724	0.215	0.652
	PCA	0.458	0.170	0.298	0.571
	BotCL	0.618	0.453	0.281	0.835
$k = 15$	ACE	0.614	0.221	0.151	—
	k-means	0.816	<b>0.978</b>	0.272	0.747
	PCA	0.432	0.162	0.286	0.645
	BotCL	<b>0.925</b>	0.744	<b>0.452</b>	<b>0.998</b>

Table 3. Results of our user study.

Dataset	Concepts	CDR $\uparrow$		CC $\uparrow$		MIC $\downarrow$	
		Mean	Std	Mean	Std	Mean	Std
MNIST	Annotated	1.000	0.000	0.838	0.150	0.071	0.047
	BotCL	0.825	0.288	0.581	0.274	0.199	0.072
	Random	0.122	0.070	0.163	0.074	0.438	0.039
CUB200	Annotated	0.949	0.115	0.595	0.113	0.512	0.034
	BotCL	0.874	0.156	0.530	0.116	0.549	0.036
	Random	0.212	0.081	0.198	0.039	0.574	0.031

of MNIST and 30 participants for CUB200 using Amazon Mechanical Turk.

We defined three metrics to summarize the participants’ responses. (i) **Concept discovery rate (CDR)**: The ratio of the responses that are not *None of them* to all responses. A higher CDR means participants can find some consistent visual elements for many concepts. (ii) **Concept consistency (CC)**: The ratio of exact matches out of all pairs of participants’ responses. A high value means many participants use the same phrases to describe a concept. (iii) **Mutual information between concepts (MIC)**: The similarity of the response distribution, computed over all possible pairs of concepts. This value is high when multiple concepts cover

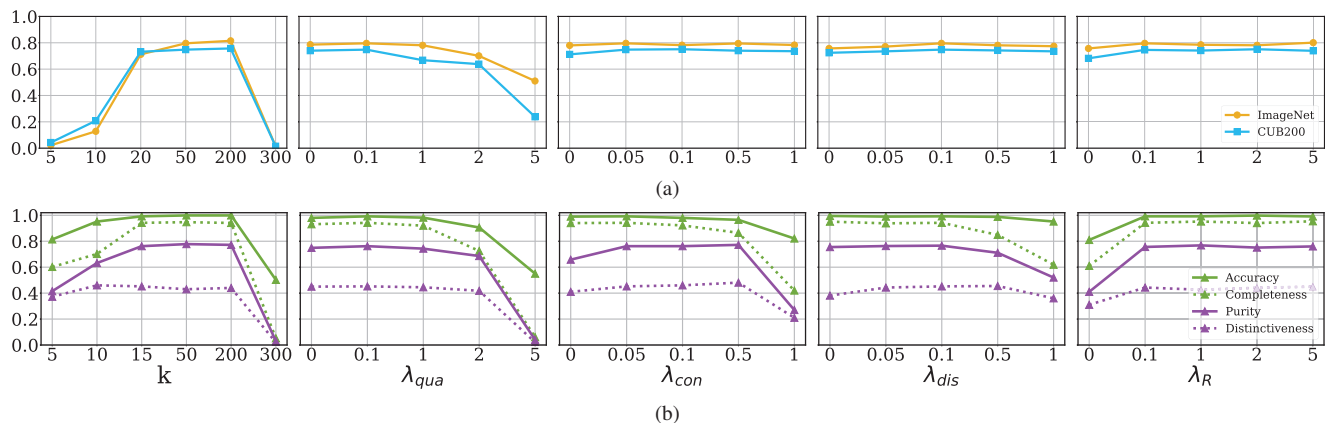


Figure 7. Results of ablation study. (a) Hyperparameter values vs. classification accuracy on ImageNet and CUB200. (b) Hyperparameter values vs. classification accuracy and other metrics on Synthetic.

the same visual elements; therefore, lower is better.

For comparison, we also evaluated a manual annotation<sup>10</sup> and random scribbling for the same images. Table 3 shows that BotCL yields good scores for all metrics on both datasets (close to the manual annotation), showing the learned concepts are interpretable for humans (from CDR), consistent (from CC), and mutually distinct (from MIC). More details are in **supp. material**.

## 5.6. Ablation Study

We conducted ablation studies using the default hyperparameters except for the one to be explored. As there is no ground truth concept for CUB200 and ImageNet, only accuracy is evaluated (Figure 7a). For Synthetic, accuracy and the three metrics in Section 5.4 are employed (Figure 7b).

**Impact of  $k$ .** A small  $k$  decreases the accuracy and other metrics, which means the necessity of searching the minimum number of concepts. Also, training tends to fail for all datasets when  $k$  is large (detailed in the **supp. materials**). The number of concepts should be tuned for each dataset. This sensitivity is one of BotCL’s limitations.

**Impact of  $\lambda_{qua}$ .** This hyperparameter controls how close  $t$  should be to a binary. The accuracy and the other metrics worsen when  $\lambda_{qua}$  gradually increases. BotCL encodes some information into  $t$  (such as the area that a concept occupies), which is lost for larger  $\lambda_{qua}$ . An extreme value may also cause vanishing gradients.

**Impact of  $\lambda_{con}$  and  $\lambda_{dis}$ .** The individual consistency and mutual distinctiveness losses hardly affect the performance on CUB200 and ImageNet, although we can see a slight drop when the values are zero for CUB200. For Synthetic, the performance metrics vary as they are designed to be. Meanwhile, the accuracy is relatively insensitive to these hyperparameters. The choice of concepts may be

highly arbitrary, and different sets of concepts may achieve similar classification performance. This arbitrariness may allow the designing of dedicated concept regularizers for the target task. However, training failures happen when they are set to be large. A small value benefits training.

**Impact of  $\lambda_R$ .** Due to the lower performance of the reconstruction loss, we studied the impact of the contrastive loss only. The contrastive loss almost always improves the classification accuracy. The performance boost is significant in CUB200 and Synthetic. As ImageNet has more training data, this may imply that self-supervision greatly contributes to the learning of concepts when training samples are limited. These results demonstrate the importance of the contrastive loss. This is interesting since this loss uses the same labels as the classification loss.

## 6. Conclusion

This paper presents BotCL for learning bottleneck concepts. Our qualitative and quantitative evaluation showed BotCL’s ability to learn concepts without explicit supervision on them but through training for a target classification task. We also demonstrated that BotCL could provide interpretability on its decision and learned concepts themselves. **Limitations.** One limitation of BotCL is that it requires tuning the number  $k$  of concepts for each dataset. It might be an interesting research direction to estimate  $k$ , *e.g.*, based on the number of classes in a given classification task. We will investigate the phenomenon to mitigate this problem.

**Acknowledgement** This work is partly supported by JST CREST Grant No. JPMJCR20D3, JST FOREST Grant No. JPMJFR2160, JSPS KAKENHI Grant-in-Aid for Scientific Research (A). This work is also supported by JSPS KAKENHI Grant Number 19K10662, 20K23343, 21K17764, and 22H03353.

<sup>10</sup>The authors annotated.



## References

- [1] David Alvarez-Melis and Tommi S Jaakkola. Towards robust interpretability with self-explaining neural networks. *NeurIPS*, 2018.
- [2] Alejandro Barredo Arrieta, Natalia Díaz-Rodríguez, Javier Del Ser, Adrien Bénénot, Siham Tabik, Alberto Barredo, Salvador García, Sergio Gil-López, Daniel Molina, and Richard Benjamins. Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Information Fusion*, 58:82–115, 2020.
- [3] Sebastian Bach, Alexander Binder, Grégoire Montavon, Frederick Klauschen, Klaus-Robert Müller, and Wojciech Samek. On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation. *PloS one*, 10(7):e0130140, 2015.
- [4] David Bau, Bolei Zhou, Aditya Khosla, Aude Oliva, and Antonio Torralba. Network dissection: Quantifying interpretability of deep visual representations. In *CVPR*, pages 6541–6549, 2017.
- [5] Maxime Bucher, Stéphane Herbin, and Frédéric Jurie. Semantic bottleneck for computer vision tasks. In *ACCV*, pages 695–712, 2018.
- [6] Aditya Chattopadhyay, Anirban Sarkar, Prantik Howlader, and Vineeth N Balasubramanian. Grad-CAM++: Generalized gradient-based visual explanations for deep convolutional networks. In *WACV*, pages 839–847, 2018.
- [7] Hila Chefer, Shir Gur, and Lior Wolf. Transformer interpretability beyond attention visualization. In *CVPR*, pages 782–791, 2021.
- [8] Chaofan Chen, Oscar Li, Daniel Tao, Alina Barnett, Cynthia Rudin, and Jonathan K Su. This looks like that: Deep learning for interpretable image recognition. *NeurIPS*, 32, 2019.
- [9] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *ICML*, pages 1597–1607, 2020.
- [10] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. In *CVPR*, pages 248–255, 2009.
- [11] Li Deng. The mnist database of handwritten digit images for machine learning research. *Signal Processing Magazine*, 29(6):141–142, 2012.
- [12] Ruth Fong, Mandela Patrick, and Andrea Vedaldi. Understanding deep networks via extremal perturbations and smooth masks. In *ICCV*, pages 2950–2958, 2019.
- [13] Yunhao Ge, Yao Xiao, Zhi Xu, Meng Zheng, Srikrishna Karanam, Terrence Chen, Laurent Itti, and Ziyang Wu. A peek into the reasoning of neural networks: Interpreting with structural visual concepts. In *CVPR*, pages 2195–2204, 2021.
- [14] Amirata Ghorbani, James Wexler, James Zou, and Been Kim. Towards automatic concept-based explanations. *NeurIPS*, 2019.
- [15] Ju He, Jie-Neng Chen, Shuai Liu, Adam Kortylewski, Cheng Yang, Yutong Bai, and Changhu Wang. Transfg: A transformer architecture for fine-grained recognition. In *AAAI*, pages 852–860, 2022.
- [16] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, pages 9729–9738, 2020.
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016.
- [18] Yusuke Hirota, Yuta Nakashima, and Noa Garcia. Quantifying societal bias amplification in image captioning. In *CVPR*, pages 13450–13459, 2022.
- [19] Andreas Holzinger, Georg Langs, Helmut Denk, Kurt Zatloukal, and Heimo Müller. Causability and explainability of artificial intelligence in medicine. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 9(4):e1312, 2019.
- [20] Zixuan Huang and Yin Li. Interpretable and accurate fine-grained recognition via region grouping. In *CVPR*, pages 8662–8672, 2020.
- [21] Been Kim, Martin Wattenberg, Justin Gilmer, Carrie Cai, James Wexler, Fernanda Viegas, and Rory Sayres. Interpretability beyond feature attribution: Quantitative testing with concept activation vectors TCAV. In *ICML*, pages 2668–2677, 2018.
- [22] Pang Wei Koh, Thao Nguyen, Yew Siang Tang, Stephen Mussmann, Emma Pierson, Been Kim, and Percy Liang. Concept bottleneck models. In *ICML*, pages 5338–5348, 2020.
- [23] Neeraj Kumar, Alexander C Berg, Peter N Belhumeur, and Shree K Nayar. Attribute and simile classifiers for face verification. In *ICCV*, pages 365–372, 2009.
- [24] Brenden M Lake, Ruslan Salakhutdinov, and Joshua B Tenenbaum. Human-level concept learning through probabilistic program induction. *Science*, 350(6266):1332–1338, 2015.
- [25] Thibault Laugel, Marie-Jeanne Lesot, Christophe Marsala, Xavier Renard, and Marcin Detyniecki. The dangers of post-hoc interpretability: Unjustified counterfactual explanations. *arXiv preprint arXiv:1907.09294*, 2019.
- [26] Liangzhi Li, Bowen Wang, Manisha Verma, Yuta Nakashima, Ryo Kawasaki, and Hajime Nagahara. Scouter: Slot attention-based classifier for explainable image recognition. In *ICCV*, pages 1046–1055, 2021.
- [27] Francesco Locatello, Dirk Weissenborn, Thomas Unterthiner, Aravindh Mahendran, Georg Heigold, Jakob Uszkoreit, Alexey Dosovitskiy, and Thomas Kipf. Object-centric learning with slot attention. *NeurIPS*, 2020.
- [28] Max Losch, Mario Fritz, and Bernt Schiele. Interpretability beyond classification output: Semantic bottleneck networks. *arXiv preprint arXiv:1907.10882*, 2019.
- [29] Vitali Petsiuk, Abir Das, and Kate Saenko. RISE: Randomized input sampling for explanation of black-box models. *BMVC*, 2018.
- [30] Andres Felipe Posada-Moreno, Nikita Surya, and Sebastian Trimpe. ECLAD: Extracting concepts with local aggregated descriptors. *arXiv preprint arXiv:2206.04531*, 2022.
- [31] Mattia Rigotti, Christoph Miksovics, Ioana Giurgiu, Thomas Gschwind, and Paolo Scotton. Attention-based interpretability with concept transformers. In *ICLR*, 2022.

- [32] Karl Schulz, Leon Sixt, Federico Tombari, and Tim Landgraf. Restricting the flow: Information bottlenecks for attribution. *arXiv preprint arXiv:2001.00396*, 2020.
- [33] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-CAM: Visual explanations from deep networks via gradient-based localization. In *CVPR*, pages 618–626, 2017.
- [34] Botian Shi, Lei Ji, Pan Lu, Zhendong Niu, and Nan Duan. Knowledge aware semantic concept expansion for image-text matching. In *IJCAI*, volume 1, page 2, 2019.
- [35] Avanti Shrikumar, Peyton Greenside, and Anshul Kundaje. Learning important features through propagating activation differences. In *ICML*, pages 3145–3153, 2017.
- [36] Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Deep inside convolutional networks: Visualising image classification models and saliency maps. In *ICLR Workshop*, 2014.
- [37] Wolfgang Stammer, Marius Memmel, Patrick Schramowski, and Kristian Kersting. Interactive disentanglement: Learning concepts by interacting with their prototype representations. In *CVPR*, pages 10317–10328, 2022.
- [38] Bas HM van der Velden, Hugo J Kuijff, Kenneth GA Gilhuijs, and Max A Viergever. Explainable artificial intelligence (XAI) in deep learning-based medical image analysis. *Medical Image Analysis*, page 102470, 2022.
- [39] Saurabh Varshneya, Antoine Ledent, Robert A Vandermeulen, Yunwen Lei, Matthias Enders, Damian Borth, and Marius Kloft. Learning interpretable concept groups in CNNs. *IJCAI*, 2021.
- [40] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, pages 5998–6008, 2017.
- [41] Bowen Wang, Liangzhi Li, Manisha Verma, Yuta Nakashima, Ryo Kawasaki, and Hajime Nagahara. Mtunet: Few-shot image classification with visual explanations. In *CVPR workshops*, pages 2294–2298, 2021.
- [42] Danding Wang, Qian Yang, Ashraf Abdul, and Brian Y Lim. Designing theory-driven user-centric explainable AI. In *Proc. CHI conference on human factors in computing systems*, pages 1–15, 2019.
- [43] Haofan Wang, Zifan Wang, Mengnan Du, Fan Yang, Zijian Zhang, Sirui Ding, Piotr Mardziel, and Xia Hu. Score-CAM: Score-weighted visual explanations for convolutional neural networks. In *CVPR workshops*, pages 24–25, 2020.
- [44] Peter Welinder, Steve Branson, Takeshi Mita, Catherine Wah, Florian Schroff, Serge Belongie, and Pietro Perona. Caltech-UCSD birds 200. 2010.
- [45] Liang Wang Yan Huang, Qi Wu. Learning semantic concepts and order for image and sentence matching. In *CVPR*, 2018.
- [46] Chih-Kuan Yeh, Been Kim, Sercan Arik, Chun-Liang Li, Tomas Pfister, and Pradeep Ravikumar. On completeness-aware concept-based explanations in deep neural networks. In *NeurIPS*, volume 33, pages 20554–20565, 2020.
- [47] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, pages 586–595, 2018.
- [48] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Object detectors emerge in deep scene CNNs. *ICLR*, 2015.
- [49] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep features for discriminative localization. In *CVPR*, pages 2921–2929, 2016.
- [50] Bolei Zhou, Yiyou Sun, David Bau, and Antonio Torralba. Interpretable basis decomposition for visual explanation. In *ECCV*, pages 119–134, 2018.