Towards Scalable and Unified Example-based Explanation and Outlier Detection

Penny Chong, Ngai-Man Cheung, Yuval Elovici, and Alexander Binder

Abstract—When neural networks are employed for high-stakes decision-making, it is desirable that they provide explanations for their prediction in order for us to understand the features that have contributed to the decision. At the same time, it is important to flag potential outliers for in-depth verification by domain experts. In this work we propose to unify two differing aspects of explainability with outlier detection. We argue for a broader adoption of prototype-based student networks capable of providing an example-based explanation for their prediction and at the same time identify regions of similarity between the predicted sample and the examples. The examples are real prototypical cases sampled from the training set via our novel iterative prototype replacement algorithm. Furthermore, we propose to use the prototype similarity scores for identifying outliers. We compare performances in terms of the classification, explanation quality, and outlier detection of our proposed network with other baselines. We show that our prototype-based networks beyond similarity kernels deliver meaningful explanations and promising outlier detection results without compromising classification accuracy.

Index Terms—prototypes, explainability, LRP, outlier detection, pruning, image classification.

I. INTRODUCTION

DEEP neural networks (DNNs) are widely used in various fields because they show superior performance. Explainable AI, aiming at transparency for predictions [1]–[3] has gained research attention in recent years. Since humans are known to generalize fast from a few examples, it is suitable to explain a prediction of a test sample via a set of similar examples. Such an example-based explanation can simply be achieved using a similarity search over an available dataset based on a metric defined by the feature maps obtained from the trained model. To achieve this, one natural option is to search in an available dataset for the nearest candidates to the test sample in the feature map space. This yields an informative visualization of the feature embedding learned by the model, yet the similar examples found do not participate in the prediction of the test sample. Even if these similar examples share the same prediction with the test sample and are close in feature space, it is not obvious to what *quantifiable extent* the model shares the same reasoning between examples and the test sample, and which parts of the test sample and

nearest neighbors share the same features. An alternative to achieve such explanation by examples is to employ kernel-based predictors [4]. These methods compute a weighted sum of similarities between the test and the training samples, and thus the impact of each training sample on the prediction of the test sample is naturally quantifiable. Although neural networks perform very well these days without the need for kernel-based setups, one may consider training a prototype-based student network from an arbitrarily structured teacher network to provide an example-based explanation¹. The prototypes in the student network participate in the network prediction, which resembles the participation of training samples in kernel-based predictions.

In this work we propose to unify two aspects of explanation with outlier detection, independent of the structure of the original model, by employing a prototype-based student network to address the three goals. The explanation for predictions is given by the network in these two aspects: i) the top- k most similar prototypical examples and ii) regions of similarity between the prediction sample and the top- k prototypes. Contrasting previous prototype learning approaches [5]–[11], we refrain from directly training prototype vectors in the latent space to avoid reconstruction errors when visualizing the learned prototype vectors in the input space. In lieu of this, we introduce an auxiliary output branch in the network for iterative prototype replacement along the lines of prototype selection methods [12]–[14] to select representative training examples for prediction inspired by dual kernel support vector machines (SVMs) [4]. Our method omits the need to map or decode prototype vectors and guarantees an explanation relative to real examples present in the training distribution. The prototype importance weight learned via the prototype replacement auxiliary loss also enables the pruning of uninformative prototypes.

With the proposed approach, it is imperative to consider the question of performance tradeoff when converting a standard CNN teacher network into a prototype-based student network. We quantify the performance of the predictors in the following three aspects: the prediction accuracy, the quality of explanation of the student network compared to the teacher network, and finally the competence of each predictor in quantifying the outlieriness of a given sample. The first two aspects are the natural tradeoffs to be considered when training any surrogate model, while the last aspect measures the outlier sensitivity of predictors in high-stakes environments. It is necessary for

P. Chong and N.-M. Cheung are with the Information Systems Technology and Design Pillar, Singapore University of Technology and Design, 487372, Singapore. Corresponding author: P. Chong. (e-mail: penny_chong@mymail.sutd.edu.sg; ngaiman_cheung@sutd.edu.sg)

Y. Elovici is with the Department of Software and Information Systems Engineering, Ben-Gurion University of the Negev, Beer-Sheva, 84105, Israel. (email: elovici@bgu.ac.il)

A. Binder is with the Department of Informatics, University of Oslo, 0315 Oslo, Norway. (email: alexabin@ifi.uio.no)

¹The teacher network is a typical convolutional neural network (CNN) for image classification tasks. The soft labels obtained from the teacher will be used for student-teacher learning to train a prototype-based student network.

predictors deployed in high-stakes settings to be equipped with the capability to flag inputs that are either anomalous or poorly represented by the training set for additional validation by experts. Prototype-based approaches deliver this property naturally by the sequence of sorted similarities between the test sample and the prototypes even though they are not trained primarily for outlier detection. The following summarizes our contribution in this work:

- (1) We argue for unified interpretable models capable of explaining and detecting outliers with prototypical examples. We demonstrate that a student network based on prototypes is capable of performing two types of explanation tasks and one outlier task simultaneously. We select prototypes from the training set that guarantee that the explanation based on examples comes from the true data distribution.
The prototype network provides two layers of explanation, first by identifying the top- k nearest prototypical examples, and second, by showing the pixel evidence of similarity between each of these examples and the test sample. The latter is achieved by backpropagating scores from similarity layers using the Layer-wise Relevance Propagation (LRP) approach for explainability.
- (2) We introduce a novel iterative prototype replacement algorithm relying on training with masks and a widely reusable auxiliary loss term. We demonstrate that this prototype replacement approach scales to larger datasets such as LSUN and PCam.
- (3) We propose various generic head architectures that compute prototype-based predictions from a generic CNN feature extractor and evaluate their prediction accuracy, explanation quality, and outlier detection performance.
- (4) We revisit a method to quantify the degree of outlierness for a sample based on the sorted prototype similarity scores. By doing so, we avoid the complex issue of hyperparameter selection in unsupervised outlier detection methods without resorting to problematic tuning of test sets. A smaller contribution lies in the derivation of the prototype similarity scores depending on the network architectures.

II. RELATED WORK

Prototype Learning One of the earliest works in prototype learning is the k -nearest neighbor (k -NN) classifier [15]. In recent years, researchers have proposed combining DNNs with prototype-based classifiers. The work [5] proposed a framework known as convolutional prototype learning (CPL), where the prototype vectors are optimized to encourage representations that are intraclass compact and interclass separable. Subsequently, interpretable prototype networks [6], [7] were proposed to provide explanations based on similar prototypes for image and sequence classification tasks. The former used a decoder, while the latter projected each prototype to its closest embedding from the training set. Another work [8] introduced a prototypical part network (ProtoPNet) that is able to present prototypical cases that are similar to the parts the network looks at. These explanations are used during classification

and not created post hoc. HPnet [9] classifies objects using hierarchically organized prototypes that are predefined in a taxonomy and is a generalization of ProtoPNet. Explanation-by-examples such as the use of prototypical examples is often preferred over superimposition explanations [1], [16], [17] for nontechnical end-users [18]. Further works in this domain include [10], [11].

In contrast to the aforementioned methods, we select prototypes directly from the training set and do not use trainable prototype vectors. Related works on prototype selection are [12], [13], which involve mining representative samples such as prototypes and criticism samples to provide an optimally compact representation of the underlying data distribution. Our work is related to ProtoAttend [14], which selects input-dependent prototypes via an attention mechanism applied between the input and the database of prototype candidates. A notable difference between our work and ProtoAttend is that our prototypes are global prototypes (fixed after network convergence), whereas they learn input-dependent prototypes that require a sufficiently large prototype candidate database and high memory consumption during training to obtain good performance [14]. In terms of explanation, our method is related to [8]–[11]. These methods explicitly focused on a patchwise representation of prototypes to deliver prototypical part explanation, whereas in our approach, a prototype represents the entire sample and we present a more fine-grained explanation in terms of pixel similarity between the prototype and test sample to explain their respective contribution to the membership of the predicted class using a post hoc explainability method [3].

Outlier Detection. We revisit the idea of quantifying outliers using statistics from the set of top- k similarity scores. For this, k -NN-based distance methods have been frequently considered in the form of summed distances [19], local outlier factors [20], [21], combinations of distance and density estimates [22], or kernel-based statistics [23], and such methods are still of recent interest [24] due to favorable theoretical properties. Unlike these works, we are interested in the performance of similarities that do not define a kernel, but rather emerge due to attention-based operations on feature maps. A simple baseline for out-of-distribution detection is the maximum softmax probability method [25] which assumes that erroneously classified samples or outliers exhibit lower probability scores for the predicted class than the correctly classified example. Several other works related to the detection of out-of-distribution samples are [26]–[28]. Generalized ODIN [27] improved upon the ODIN [26] method, which is based on temperature scaling and adding small perturbations, by proposing to decompose confidence scores and a modified input preprocessing method that eliminates the need to tune hyperparameters with out-of-distribution samples. Our formulation of the outlier score using the sum of prototype similarity scores is based on the idea in [23] and is also the generalization of the confidence score in ProtoAttend [14].

III. METHODOLOGY

In a practical development task, many researchers prefer to start with a lower complexity baseline. As of 2021, this

baseline for image classification would still be the conventional CNN for many practitioners because many people have some experience in training CNNs. Once one has trained a baseline, one has a teacher available to obtain soft labels for free. Furthermore, it is anecdotally reported that training with soft labels outperforms training with hard labels [29]; thus, we feel it is natural to adopt teacher-student training in this work. We employ teacher-student learning with soft labels for our proposed prototype-based network to transfer the knowledge from a standard teacher CNN to a prototype-based student network to retain the predictive performance of the teacher. In this section, we first introduce our prototype network architectures, the iterative prototype replacement algorithm, and the learning objectives. We also present the explainability methods for the network and the formulation of the outlier score with respect to the student architectures.

Fig. 1. Taxonomy of the different categories of head architectures proposed for the student network.

Our proposed architecture for the student network is inspired by the dual formulation of SVMs [4] that solves the optimization problem $\max_{\alpha_i \geq 0} \sum_i \alpha_i - \frac{1}{2} \sum_j \sum_k \alpha_j \alpha_k y_j y_k K(\mathbf{x}_j, \mathbf{x}_k)$, subject to the constraints $0 \leq \alpha_i \leq C$, $\forall i$ and $\sum_i \alpha_i y_i = 0$. For a binary problem, the dual classifier takes the form of $f(\mathbf{x}) = \sum_i \alpha_i y_i K(\mathbf{x}_i, \mathbf{x}) + b$, where $K(\mathbf{x}_i, \mathbf{x})$ is a positive definite kernel. The solution is based on the set of similarities $\{K(\mathbf{x}_i, \mathbf{x}) \mid 1 \leq i \leq N\}$ between the training samples and the test sample $\mathbf{x}$ which can be used to compute statistics to quantify outliers [23]. We employ architectures based on attention-weighted similarities beyond Mercer-Kernels.

A. Proposed Architecture for the Student Network

For the purpose of providing meaningful explanations and detecting outliers, we introduce three generic head architectures² for our prototype student network. Each architecture is evaluated for its performance in providing explanation and detecting outliers in Section V. For simplicity, we assume in this section that the prototypical examples are already selected based on our iterative prototype replacement algorithm. We will introduce the algorithm in Section III-B later.

Let $f(\cdot; \theta_1) : \mathbb{R}^{C_0 \times H_0 \times W_0} \rightarrow \mathbb{R}^{C \times H \times W}$ be the CNN feature extractor before the spatial pooling layer with the learnable parameters θ_1 . Let $e(\cdot; \theta_2) : \mathbb{R}^{C \times H \times W} \rightarrow \mathbb{R}^{\tilde{C}}$ be

²The term head architecture refers to the architecture used in the prototype network which is the portion of the network located after the CNN encoder. Figures 2-4 illustrate this.

the prototype network with the learnable parameters θ_2 , and let $\tilde{C}$ be the number of classes in the classification task. The overall student network is represented by $e \circ f$. For simplicity, we omit the parameters θ_1 and θ_2 and denote the feature map representation of the input sample x as $f(x)$ in this section.

To make a prediction, we have to solve the problem of computing a similarity between the feature map $f(x)$ for a test sample x and the feature $f(p)$ for a prototype p . The feature maps $f_{c,h,w}(\cdot)$ have a channel dimension c and spatial dimensions, usually the two h, w . We observe that three larger categories of similarities can be defined, which correspond to the head classes **I**, **II** and **III**. The **Head I** follows the simplest idea to compute a similarity without any spatial dimensions; thus, it pools over the spatial dimensions before computing an inner product over the channel dimension. The **Head II** class computes similarities that incorporate the spatial dimensions. The difference between the two heads in the Head II class is whether one matches each spatial coordinate (h, w) in the test sample to the same coordinate in the prototype as in Head II-A or whether one searches for every coordinate in the test sample for the best matching coordinate in the prototype using a maximum-operation as in Head II-B. Finally, the **Head III** class also uses the spatial dimensions of feature maps but additionally learns an attention weight over the spatial coordinates of the features in the test sample. This serves to weight the spatial regions in the test samples in the final similarity measures, whereas the **Head II** class treats each spatial coordinate (h, w) of the test sample feature map with equal weight.

Figure 1 shows the taxonomy of the various head architectures for the student network and Figures 2-4 are visualizations of the head architectures. Our task differs from the standard image retrieval task as our task uses prototypes from the training set that participate in the prediction and thus share quantifiable prediction reasoning as the test sample, whereas the image retrieval task involves an exhaustive nearest neighbor search over the entire database that does not participate in the prediction and thus does not guarantee similar prediction reasoning between the test sample and its nearest neighbor.

Our proposed head architectures do not require any modification to the CNN encoder compared to [14], which requires additional linear layers for queries and keys that add to the overhead during training.

The **Head I** architecture is based on a spatial average pooling of the feature maps $f(\cdot)$, resulting in a vector $g(\cdot)$ that consists of only C feature channels. This is followed by a cosine similarity over the feature channels $g(\cdot)$, as seen in Figure 2:

$$s^{(I)}(p_k) = \frac{g(x)^\top g(p_k)}{\|g(x)\| \|g(p_k)\|}. \quad (1)$$

The similarity output $s^{(I)}(p_k)$ falls in the range of $[-1, 1]$ where a higher value indicates higher similarity between the input sample x and the prototype p_k . The linear classification layer $h(x, \{p_1, p_2, \dots, p_K\}) = \sum_{k=1}^K w_k \text{ReLU}(s^{(I)}(p_k)) + b$ then predicts the output logit y . For this head architecture, the learnable prototype network parameters θ_2 denote the parameters of the linear layer $h(\cdot)$. The linear layer $h(\cdot)$ in

Fig. 2. Head I : Similarity using features after spatial pooling. The y_{mask} output branch is the auxiliary output branch for iterative prototype replacement.

Fig. 3. Head II : Similarity using features before spatial pooling. The y_{mask} output branch is the auxiliary output branch for iterative prototype replacement.

Fig. 4. Head III : Similarity w/attention using features before spatial pooling. The y_{mask} output branch is the auxiliary output branch for iterative prototype replacement.

Head I is consistent with the formulation of the dual kernel SVM. The linear layers $h(\cdot)$ in Heads II and III as shown in Figures 3 and 4 will have similar dependency on a set of prototypes $\{p_1, p_2, \dots, p_K\}$, but they no longer correspond to kernels.

The **Head II** architecture in Figure 3, uses the features before spatial pooling to compute similarity scores. Therefore, the similarity layer computes $s^{(II)}(p_k) \in \mathbb{R}^{H \times W}$, i.e., one similarity score for each spatial position (h, w) . We explore two types of similarity functions for this head architecture: A) similarity between the input and prototype at the same spatial position (h, w) and B) the maximum similarity between input and prototype at a given *input spatial location* (h, w) . We refer to the former as Head II-A and the latter as Head II-B. The similarity operation for Head II-A (Eq.(2)) and Head II-B (Eq.(3)) at the spatial location (h, w) is defined as:

$$s_{h,w}^{(II-A)}(p_k) = \sum_c \hat{f}_{c,h,w}(x) \cdot \hat{f}_{c,h,w}(p_k), \quad (2)$$

$$s_{h,w}^{(II-B)}(p_k) = \max_{h',w'} \sum_c \hat{f}_{c,h,w}(x) \cdot \hat{f}_{c,h',w'}(p_k), \quad (3)$$

$$\hat{f}_{c,h,w}(\cdot) = \frac{f_{c,h,w}(\cdot)}{\|f_{c,h,w}(\cdot)\|_2} = \frac{f_{c,h,w}(\cdot)}{\sqrt{\sum_c [f_{c,h,w}(\cdot)]^2}}, \quad (4)$$

where $f_{c,h,w}(x)$ and $f_{c,h,w}(p_k)$ are the c -th features at the spatial location (h, w) from the feature maps $f(x)$ and $f(p_k)$, respectively. Head II-A compares the pairwise similarity of the spatial features between $f(x)$ and $f(p_k)$ along the same spatial location, whereas Head II-B considers all spatial locations of $f(p_k)$ to find the location that is most similar to the specified spatial location in $f(x)$. Therefore, Eq. (3) ensures that the similarity at the spatial location (h, w) is computed only between the location (h, w) in x and the most relevant location (h', w') in p_k , i.e., the pair that locally gives the highest

cosine similarity score. Its similarity operation is spatially invariant, as it considers all spatial locations of prototype p_k in identifying the most similar feature to the input x . Head II-A (Eq. (2)) can be viewed as a special case of Head II-B (Eq. (3)). The operation of Head II-B reduces to Head II-A when the maximum similarity between input and prototype for a given *input spatial location* (h, w) , also occurs at the prototype location (h, w) . Thus, the Head II-B architecture is considered the general case of Head II-A. For this category of architectures, the learnable prototype network parameters θ_2 denote the parameters of the linear layer $h(\cdot)$. Note that Head II-B, as a maximum over inner products, no longer defines a kernel, as the maximum of two positive definite matrices is not guaranteed to be a positive definite. Thus, we employ a nonkernel-based similarity.

The **Head III** architecture employs a common attention mechanism to compute weighted similarities. For the Head III architecture as shown in Figure 4, the similarity scores from Eq. (2) and (3) are passed into a softmax $\sigma(\cdot)$ to obtain the following attention maps:

$$a_{h,w}^{(III-A)}(p_k) = \sigma(s_{h,w}^{(II-A)}(p_k)), \quad (5)$$

$$a_{h,w}^{(III-B)}(p_k) = \sigma(s_{h,w}^{(II-B)}(p_k)). \quad (6)$$

Head III-A and Head III-B are the attention-augmented counterparts of Head II-A and Head II-B, respectively:

$$s_c^{(III-A)}(p_k) = \sum_{h,w} a_{h,w}^{(III-A)}(p_k) \cdot f_{c,h,w}(x) \cdot f_{c,h,w}(p_k), \quad (7)$$

$$s_c^{(III-B)}(p_k) = \sum_{h,w} a_{h,w}^{(III-B)}(p_k) \cdot f_{c,h,w}(x) \cdot f_{c,h_{\max},w_{\max}}(p_k). \quad (8)$$

The indices $h_{\max}$ and $w_{\max}$ in Eq. (8) are the selected h' and w' indices used in the computation of $s_{h,w}^{(II-B)}(p_k)$. The attention

Algorithm 1 Iterative Prototype Replacement Algorithm

```

1: Initialization: Let  $S = D \cup P$  be the original training set with  $n(S) = n(D) + n(P) = N + K$ .  $D$  as the new training
   set and  $P = \{p_1, \dots, p_k, \dots, p_K\}$  as the set of all prototypes. The initial set  $P$  consists of an equal number of samples per
   class that are randomly sampled from the original training set  $S$ . Let  $\theta_1$  and  $\theta_2$  be the parameters of the CNN feature
   extractor and the prototype network, respectively, and  $\mathcal{M} \in \mathbb{R}^K$  be the prototype importance weight used for replacement.
2: while  $e \leq$  number of epochs do
3:   while  $t \leq$  number of iterations do
4:     Perform a forward pass to obtain input features  $f(x)$ , set of prototype features  $\{f(p_k)\}$ , corresponding input features
        $\{z(p_k)\}$  of the linear classification layer  $h(\cdot)$ , and the output logits  $y$ . (For Head I, we compute  $g(x)$  and  $\{g(p_k)\}$ 
       instead of  $f(x)$  and  $\{f(p_k)\}$ .)
5:     Compute threshold  $\tau_t$  as the  $p$ -th smallest weight in  $\mathcal{M}$ .
6:     Compute binary prototype mask  $\mathcal{M}'(\tau_t)$  from  $\mathcal{M}$  as shown in Eq. (11).
7:     Perform a forward pass in  $h(\cdot)$  using the masked input to obtain the auxiliary output  $y_{mask}$  as shown in Eq. (12).
8:     Optimize parameters  $\theta_1, \theta_2$  and  $\mathcal{M}$  based on the loss function (13).
9:   end while
10:   $P \leftarrow P \setminus P_{replace} \cup D_{rand}$ , where  $P_{replace}$  is the set of  $p$  prototypes with the smallest weight in  $\mathcal{M}$ , and  $D_{rand}$  is the
    set of random samples comprising the replaced prototype classes sampled from  $D$  to replace  $P_{replace}$ .
11:   $D \leftarrow D \setminus D_{rand} \cup P_{replace}$ 
12:  Reinitialize only the weights in  $\mathcal{M}$  where prototypes are replaced.
13: end while

```

map $a(p_k) \in \mathbb{R}^{H \times W}$ is applied over the spatial features $f(x) \in \mathbb{R}^{C \times H \times W}$ and $f(p_k) \in \mathbb{R}^{C \times H \times W}$ to give more weight to the important regions (i.e., spatial regions between $f(x)$ and $f(p_k)$ with high similarity) for the neural network to learn and focus on these regions. The output from the attended similarity layer which represents a weighted spatial similarity for every channel, is passed into a convolutional-1D layer with kernel size C to compute a weighted sum of the C features. Since $C \gg H \times W$, we avoid using simple average spatial pooling over the C features to obtain better explanation heatmaps. The convolutional-1D layer has a minimum weight clipping at zero to ensure the outputs are always nonnegative in order to apply the explainability method described in Section III-D. Additionally, the use of individual attention maps for the respective features $f(x)$ and $f(p_k)$ is also investigated in Head III-C:

$$a_{h,w}^{(III-C)}(p_k) = \sigma \left(\max_{u,v} \sum_c \hat{f}_{c,u,v}(x) \cdot \hat{f}_{c,h,w}(p_k) \right), \quad (9)$$

$$s_c^{(III-C)}(p_k) = \sum_{h,w} a_{h,w}^{(III-B)}(p_k) \cdot f_{c,h,w}(x) \cdot a_{h,w}^{(III-C)}(p_k) \cdot f_{c,h,w}(p_k), \quad (10)$$

where $a_{h,w}^{(III-B)}$ is the attention score from Eq. (6) for $f(x)$ at the input spatial location (h, w) , whereas $a_{h,w}^{(III-C)}$ is the attention score for $f(p_k)$ at the prototype spatial location (h, w) , as can be seen by the computation of $\max_{u,v}(\cdot)$ over the spatial locations of $\hat{f}(x)$. Eq. (9) computes an attention map by considering all spatial locations of $f(x)$ to find the most similar location to the specified location in $f(p_k)$ as opposed to the attention map based on Eq. (3). The attention map $a^{(III-B)}(p_k) \in \mathbb{R}^{H \times W}$ with weights denoting the importance of the spatial positions in $f(x)$ based on its similarity with $f(p_k)$, is applied over the feature $f(x)$. On the other hand, the attention map $a^{(III-C)}(p_k) \in \mathbb{R}^{H \times W}$ with importance weights for the spatial positions in $f(p_k)$ based on its similarity

with $f(x)$, is applied over the feature $f(p_k)$. For the head architectures in this category, the learnable prototype network parameters θ_2 are the parameters of the convolutional-1D and the linear classification $h(\cdot)$ layers.

B. Auxiliary Output Branch for Iterative Prototype Replacement

For the selection of representative prototypes, one can directly learn an importance mask over the entire training set and select those samples with high importance as prototypes. This approach, however, does not scale well to large sample sizes as the number of parameters that must be learned increases proportionally. For approaches such as [14], the database of prototype candidates used must be sufficiently large for successful training as reported in their paper. Instead of learning an importance mask over the entire training set or a large candidate database, we propose to learn a smaller importance mask over a fixed number of K prototypes. We note that K is very much smaller than the size of the candidate database [14] used during training or inference. Our approach scales well to large sample sizes since the number of parameters to be learned can now be constrained by K which also corresponds to the final number of prototypes used during inference. We introduce an auxiliary output branch to guide the network to learn a prototype replacement scheme that replaces uninformative prototypes (out of the K prototypes) with new samples based on the importance or contribution of the current prototypes to the decision boundary of the network. Prototypes that are assigned larger importance weights have higher importance and are unlikely to be replaced. On the other hand, prototypes with smaller importance weights are likely to be replaced with other samples from the training set due to their minimal contribution to the decision boundary of the network.

During training, the network learns a prototype importance weight $\mathcal{M} \in \mathbb{R}^K$ via an auxiliary loss to quantify the

importance of the K prototypes based on their contribution to the prediction. The auxiliary loss imposes a penalty when informative prototypes that influence the decision boundary of the network are removed, i.e., masked out. The network parameters θ_1 and θ_2 and the prototype importance weight $\mathcal{M}$ are jointly optimized. The formulation of the auxiliary loss will be elaborated further in Section III-C.

The full algorithm for the iterative prototype replacement method is presented in Algorithm 1. We compute the binary prototype mask $\mathcal{M}'(\tau_t)$ from the prototype importance weight $\mathcal{M}$ using the threshold τ_t at iteration t as:

$$\mathcal{M}'(\tau_t) = \frac{\max(0, \mathcal{M} - \tau_t)}{|\mathcal{M} - \tau_t|} \in \{0, 1\}. \quad (11)$$

We use continuous ReLU activation with normalization as inspired by the hard shrinkage operation in [30] to avoid issues with the backpropagation of a discontinuous 0-1 step function. For simplicity, we assign τ_t as the p -th smallest weight in $\mathcal{M}$ to zero out a fixed number of p prototypes (in the input of the linear classification layer $h(\cdot)$) with the smallest weight in $\mathcal{M}$. To compute the auxiliary output y_{mask} , we multiply the binary prototype mask $\mathcal{M}'(\tau_t)$ elementwise with the features $z(p_k) \in \mathbb{R}$ at the input of the final linear classification layer $h(\cdot)$ corresponding to each prototype as follows:

$$y_{mask} = h(x, \{p_1, p_2, \dots, p_K\}) = \sum_{k=1}^K w_k m'_k z(p_k) + b, \quad (12)$$

where $m'_k \in \{0, 1\}$ is the weight in $\mathcal{M}'(\tau_t)$, and $w_k \in \mathbb{R}$ and $b \in \mathbb{R}$ are the parameters of the linear layer $h(\cdot)$. At the end of each epoch, we replace the p prototypes that have the smallest weight in $\mathcal{M}$ with new samples from the training set D . The new set of prototypes P also maintains an equal number of samples per class during training.

The prototype importance weight $\mathcal{M}$ can also be used to prune prototypes and relevant neurons. This will be discussed in Section VI-B.

We remark that it is not of interest to iterate through all the training samples when choosing optimal prototypes, as this will not scale well to large datasets. Our algorithm is based on the idea that one can find prototypes that may be similar to a number of training samples and thus can avoid iterating through all training samples when selecting prototypes. In future work, we will consider presorting the prototypes based on their similarity to maintain an order over training samples during learning to speed up model convergence. The prototype-based Head III architectures with higher architectural complexities than Head I and II architectures have greater training overhead than a standard CNN teacher network due to the computation of attention maps and training time which increases proportionally with the number of prototypes used. Additional details are provided in Section S-I of the Supplementary Material.

C. Learning Objective

We define the set of input training samples with their ground truth to be $D = \{(x_i, \tilde{y}_i) \mid 1 \leq i \leq N\}$ and the set of prototypes with their class labels to be $P = \{(p_k, \tilde{q}_k) \mid 1 \leq$

$k \leq K\}$. For an input sample x , we denote the output logits of the teacher network as y_T and the logits and auxiliary output of the student network as y and y_{mask} , respectively. The predicted class label from the student network for the input sample x is represented by $y_{pred} = \arg \max y$. The following defines the objective function for our student network:

$$\min_{\theta_1, \theta_2, \mathcal{M}} \mathcal{L}(\tilde{y}, \sigma(y)) + \lambda_1 \mathcal{L}(\sigma(y_T), \sigma(y)) + \lambda_2 \mathcal{L}(y_{pred}, \sigma(y_{mask})) + \lambda_3 \mathcal{J}((x, \tilde{y}), (p, \tilde{q})), \quad (13)$$

where $\mathcal{L}(\cdot)$ is the cross-entropy loss, $\mathcal{J}(\cdot)$ is a loss term based on the squared L_2 -norm, and $\sigma(\cdot)$ is the softmax.

The first two terms are the losses commonly employed in the training of a student network. The second term, also known as the distillation loss, is the cross-entropy loss with the teacher soft labels as ground truth to encourage the student network to learn predictions resembling the teacher network.

The third term $\mathcal{L}(y_{pred}, \sigma(y_{mask}))$ is used for the learning of prototype importance weight $\mathcal{M} \in \mathbb{R}^K$. We use the standard cross-entropy loss as the auxiliary loss but with the predicted class label and the probability $\sigma(y_{mask})$ as the input argument to the standard cross-entropy loss rather than using the ground truth class $\tilde{y}$ and the probability output $\sigma(y)$ from the main network branch. The auxiliary loss enables the network to learn prototype importance weight $\mathcal{M} \in \mathbb{R}^K$ which gives larger weights to prototypes that will cause a large change in the decision boundary of the network when removed, and lower importance weights to prototypes that cause only a small change to the decision boundary if removed. Our proposed solution is based on the idea of removing uninformative prototypes with minimal disruption to the decision boundary using the predicted class label y_{pred} instead of the ground truth class label $\tilde{y}$ to guide the learning of parameters $\mathcal{M} \in \mathbb{R}^K$ since we believe the network can learn better by replacing uninformative prototypes with other potentially more informative prototypes. This allows the replacement of lower ranked prototypes that have a lower impact on the prediction accuracy in the prototype replacement step.

The fourth loss term $\mathcal{J}((x, \tilde{y}), (p, \tilde{q}))$ encourages the feature space of the input sample x to be close to the set of prototypes from the same class but farther away from the other classes. The formulation of this objective term varies slightly between the different head architectures depending on their similarity operation.

For the Head I architecture, the objective term $\mathcal{J}((x, \tilde{y}), (p, \tilde{q}))$ is defined as:

$$\frac{1}{|D||P|} \sum_{i=1}^{|D|} \sum_{k=1}^{|P|} (\|\hat{g}(x_i) - \hat{g}(p_k)\|_2^2)^\alpha, \quad (14)$$

where $\alpha = 1$, if $\tilde{y}_i = \tilde{q}_k$, else $\alpha = -1$.

For the Head II-A and Head III-A architectures, the objective function $\mathcal{J}((x, \tilde{y}), (p, \tilde{q}))$ is defined as:

$$\frac{1}{|D||P|} \sum_{i=1}^{|D|} \sum_{k=1}^{|P|} \left[\frac{1}{HW} \sum_{h,w,c}^{H,W,C} (\hat{f}_{c,h,w}(x_i) - \hat{f}_{c,h,w}(p_k))^2 \right]^\alpha, \quad (15)$$

where $\alpha = 1$, if $\tilde{y}_i = \tilde{q}_k$, else $\alpha = -1$.

The term inside the square bracket is simply the average squared L_2 -norm computed for the channel dimension over all spatial positions (h, w) . Similarly, the objective function $\mathcal{J}((x, \tilde{y}), (p, \tilde{q}))$ for the Head II-B and Head III-B architectures is defined as:

$$\frac{1}{|D||P|} \sum_{i=1}^{|D|} \sum_{k=1}^{|P|} \left[\frac{1}{HW} \sum_{h,w,c}^{H,W,C} (\hat{f}_{c,h,w}(x_i) - \hat{f}_{c,h_{\max},w_{\max}}(p_k))^2 \right]^\alpha, \quad (16)$$

where $\alpha = 1$, if $\tilde{y}_i = \tilde{q}_k$, else $\alpha = -1$.

The different formulation for Eq. (15) and Eq. (16) is due to the two different types of similarity operations introduced in the head architectures, as discussed in Section III-A. For Head III-C architecture, the objective function $\mathcal{J}(\cdot)$ is expressed as

$$\mathcal{J}((x, \tilde{y}), (p, \tilde{q})) = \mathcal{J}_x((x, \tilde{y}), (p, \tilde{q})) + \mathcal{J}_p((x, \tilde{y}), (p, \tilde{q})), \quad (17)$$

where the first term is equivalent to Eq. (16) and the second term is defined accordingly based on the $a_{h,w}^{(\text{III-C})}(p_k)$ operation for $f(p_k)$ in Head III-C, i.e., use $\hat{f}_{c,h_{\max},w_{\max}}(x_i)$ instead of $\hat{f}_{c,h,w}(x_i)$ and vice versa for p_k in the loss $\mathcal{J}_p(\cdot)$.

D. Explanation for Top- k Nearest Prototypical Examples

The student network provides a prediction explanation in the form of: (1) a prediction explanation with top- k nearest prototypical examples and (2) pixel evidence of similarity between the prediction sample and the prototypical examples. We define a prototype similarity score u_k that ranges between $[0, 1]$ to identify the top- k nearest prototypical examples for a given input sample. We elaborate on the formulation of the prototype similarity score u_k in Section III-E as the scores are also used to formulate the outlier score. In the following paragraph, we recapitulate the explanation method to compute pixel evidence of similarity between each pair of input-prototype.

To show evidence of similarity between the input sample and the prototypes in the pixel space, we use a posthoc explainability method, the Layer-wise Relevance Propagation (LRP) algorithm [3], [31] to compute a pair of LRP heatmaps for each pair of input-prototype. For a pair of input-prototype, their respective heatmaps comprise positive and negative relevance scores (also known as evidence) in the pixel space. The relevance scores represent the pixels' contribution to the prediction score. For the 2D-convolutional layers, we used the following LRP- $\alpha\beta$ rule to compute relevance $R_d^{(l)}$ for the neuron d at the current layer l :

$$R_d^{(l)} = \sum_j \left(\alpha \frac{(a_d w_{dj})^+}{\sum_{0,d} (a_d w_{dj})^+} - \beta \frac{(a_d w_{dj})^-}{\sum_{0,d} (a_d w_{dj})^-} \right) R_j^{(l+1)}, \quad (18)$$

where a_d is the lower-layer activation from the input neuron d at layer l , w_{dj} is the weight connecting the neuron d at layer l to the higher-layer neuron j at the succeeding layer $l+1$, $R_j^{(l+1)}$ is the relevance for the neuron j at the succeeding layer $l+1$, $\alpha - \beta = 1$, $\alpha > 0$, and $\beta \geq 0$. For the other layers,

including the similarity layer (Heads I and II) and the attended similarity layer (Head III), we use the following LRP- ϵ rule:

$$R_d^{(l)} = \sum_j \frac{a_d w_{dj}}{\epsilon + \sum_{0,d} a_d w_{dj}} R_j^{(l+1)}. \quad (19)$$

The similarity/attended similarity layer is treated as a linear layer, which enables the easy application of the LRP- ϵ rule in the standard manner without resorting to more complex LRP methods such as BiLRP [32]. For Head III architectures, the relevance scores flow through the attended similarity layer branch only while treating the attention weights from the similarity layer as constants.

From the last linear layer up to the similarity/attended similarity layer l , we compute the relevance scores $R_{(\cdot)}^{(l)}(x, p_k)$ for each pair of (x, p_k) . Using $R_{(\cdot)}^{(l)}(x, p_k)$ at the similarity/attended similarity layer l , we compute the LRP heatmap for x by backpropagating the relevances to the input space of x through the CNN encoder. The heatmap for p_k is also computed in a similar manner (for Head II-B and Head III-B architectures with $\max_{h',w'}(\cdot)$ operation, the relevance scores are accumulated at the relevant h' and w' indices before backpropagating to the input space of p_k). We refer the reader to [3] for more information on LRP algorithms.

E. Quantifying the Degree of Outlierness of a Sample

Our proposed prototype-based architectures can naturally quantify the degree of outlierness of an input sample based on the sequence of sorted similarities even though they are not trained primarily for outlier detection. We revisit the method in [23] and generalize the confidence score in [14] to compute an outlier score based on the sum of similarity scores. We use the sum of the prototype similarity scores u_k to compute the outlier score. While we see the main contribution in unifying two aspects of explanation and outlier detection, a smaller contribution lies in the different formulation of the prototype similarity scores u_k with respect to the student architecture (refer to the output branch for outlier score in Figures 2-4).

For each input sample x , we define the set of corresponding prototype similarity scores as $U = \{u_k \mid 1 \leq k \leq K\}$. For all Head I and Head II architectures, the u_k score is assigned as $u_k = z(p_k)$, where $z(p_k)$ is the input to the linear layer $h(\cdot)$ as shown in Figures 2 and 3. For Head III-A and Head III-B architectures, the u_k score is computed as $u_k = \frac{1}{HW} \sum_{h=1, w=1}^{H,W} r_{h,w}(p_k)$. For the Head III-C architecture, the u_k score is computed as $u_k = \max_{h,w}(r_{h,w}(p_k))$. With the placement of ReLU layers in the CNN, the elements of U lie in the range of $[0, 1]$, where a higher value indicates a larger similarity between the input sample x and the prototype p_k . From the set U , top- k' highest scores are selected regardless of the prototype class. For sample x , the outlier score $o(x)$ is defined as:

$$o = 1 - \frac{1}{k'} \sum_{k=1}^{k'} u_k. \quad (20)$$

Note that u_k is in general no inner product, and thus, we cannot compute true metric distances and hence follow the idea in [23]. In this formulation, we assume that the prototype

similarity scores u_k for outliers are smaller than the normal samples. The above formulation is the general case, which considers prototypes from all classes in the summation, unlike the formulation of the confidence score in [14], which considers only the set of prototypes from the predicted class. We show only the results using Eq. (20), which gives better performance, and omit results using only prototypes from the same predicted class due to lack of space. Larger outlier score o indicates a larger anomaly.

IV. EXPERIMENTAL SETTINGS

Dataset. We used the LSUN [33] dataset consisting of 10 classes with an image size of 128×128 pixels, the PCam [34], [35] dataset consisting of 2 classes with an image size of 96×96 pixels, and the Stanford Cars [36] dataset with image size 224×224 pixels, where we consider only samples in the first 50 cars classes instead of the total 196 classes due to time constraints. For the LSUN dataset, we subsampled 10,000 samples per class from the given training set to save training and computational time. We used the 80 : 20 split for training and validation. The original validation split was used as the test set. For the Stanford Cars dataset, we used fewer number of training samples because we split 10% of the given training set for each class to be our validation set, whereas the setup in [8] used all of the class samples in the given training set to train the model. We performed offline data augmentation following [8] on our training set and used the augmented training set for training. Since each sample in the training set is augmented 29 times, using 10% fewer base training samples will have a larger impact on the model performance, as the complete training set used for training consists of both augmented and base samples. In our setup, we used the performance on the validation set as an early training stopping criterion and to select the best model as opposed to using the training accuracy, cluster cost, and separation cost to stop training as done in [8]. Performance for the Stanford Cars dataset is reported on its official test set using only samples from the first 50 classes. For the PCam dataset, we used the given dataset splits as they are. We refrained from the common usage of CIFAR-10, MNIST, or Fashion MNIST classes, which are known to cluster very easily and offer little potential for aggregating spatial similarities due to low spatial resolutions. These datasets do not provide an accurate performance benchmark which we explain later in Section V-A.

LRP Perturbation. We subsampled only 100 samples per class from our LSUN test set and 500 samples per class from our PCam test set to reduce computational time. We set $\alpha = 1.7$, $\beta = 0.7$, and $\epsilon = 1e-3$ for the LRP- $\alpha\beta$ and LRP- ϵ rules, respectively.

Outlier Setup. Three different types of outlier setups were used in our experiments. Setup A consists of the LSUN test set, which is labeled the normal sample set, and the test set from the Oxford Flowers 102 [37] dataset, which is labeled the outlier sample set. For Setup B, we created a synthetic outlier counterpart for each test sample in our test set (LSUN/PCam/Stanford Cars) by drawing random *strokes* of

thickness M on the original image [38]. In Setup C, we created an outlier counterpart for each test sample by manipulating the color of the image, by increasing the minimum saturation and value components, and then randomly rotating the hue component. The manipulated image has abnormal colors. Due to the large number of test samples in the PCam dataset, 1500 samples per class were sampled from the test set for the outlier detection experiments.

We showed a comparison with the common baseline method based on the maximum probability class [25], the classic isolation forest (IF) [39], generalized ODIN (GODIN) [27], and Outlier Exposure (OE) [28] to evaluate the potential of our prototype networks in distinguishing outliers from the normal sample distribution. The outlier score for the baseline method [25] is the negative probability of the predicted class such that a sample with a smaller predicted class probability has a higher outlier score. For the training of GODIN layers, i.e., $g(x)$ and $h_i(x)$, we used a learning rate of $1e-3$, and we either fixed or fine-tuned the teacher network layers using a learning rate of $1e-4$ with an early stopping training criterion using the performance on the in-distribution validation set. For the OE method, we fine-tuned the teacher network with a learning rate of $1e-3$ on one type of outlier and tested it on the other types of outliers. For fine-tuning with the out-of-distribution Oxford Flowers 102 dataset, we used the entire dataset comprising of training, validation, and test samples for training. For the hyperparameters of IF and other hyperparameters of GODIN and OE that are not stated here, we followed the hyperparameters used by the authors. We evaluated the outlier performance of the models using the area under the receiver operating characteristic (AUROC) curve, false positive rate at 95% true positive rate (FPR95), and the area under the precision-recall (AUPR) curve metrics. We reported in Section S-II of the Supplementary Material both AUPRout and AUPRin metrics, where the former treats the outlier class as the positive class, while the latter treats the inlier class as the positive class.

Model and Hyperparameter. We employed ResNet-50 and ResNet-34 architectures [40] pretrained on ImageNet as the CNN encoder. The former is used for the LSUN and PCam datasets, and the latter is used for the Stanford Cars dataset. We used the SGD optimizer with momentum 0.9 and weight decay of $1e-4$. For all head architecture networks, except Head III-B trained on the Stanford Cars dataset, the learning rates of $1e-3$ and $1e-4$ were used in the prototype network and the CNN encoder, respectively. A decay learning rate with a step size of 10 epochs and $\gamma = 0.1$ was also adopted. To reduce overfitting for the Head III-B network on the Stanford Cars dataset, we used smaller learning rates of $5e-5$ and $1e-5$ for the prototype network and the CNN encoder, respectively, with learning rate decay every 20 epochs and $\gamma = 0.2$. For the learning objectives, we set $\lambda_1 = \lambda_2 = 1$ and $\lambda_3 = 0.1$. For the iterative prototype replacement algorithm, we set $p = 30\%$ of the number of prototypes (Line 5 of Algorithm 1). We used PyTorch [41] for implementation.

We compared our methods with the teacher network, ProtoDNN [6], ProtoPNet [8], and the χ^2 -kernel SVM [42]. To control architectural effects, we used the same ResNet

architecture pretrained on ImageNet as the CNN encoder for the aforementioned methods. We followed the hyperparameters used by the authors except for the learning rate of the encoder in ProtoDNN, which we set to a learning rate of $1e-5$. The χ^2 -kernel SVM used histogram CNN features pretrained on ImageNet. Since kernel SVMs do not scale to large datasets, we used a χ^2 approximation kernel [42] and the SGD classifier with hinge loss at an optimal learning rate and tolerance of $1e-4$. We performed a grid search on the validation set with the hyperparameters $\{1e-5, 1e-4, \dots, 1\}$ to determine the regularization constant in the SGD classifier. We used 10 prototypes per class in the LSUN and Stanford Cars experiments, and 20 prototypes per class in the PCam experiments.

To ensure reproducibility, we provide the instructions to prepare the datasets and splits used in our experiments including the codes to create the synthetic outliers in GitHub³.

V. EXPERIMENTAL RESULTS

A. Prediction Performance

We begin the comparison with the predictive aspect of our proposed architectures and the other methods on both datasets. Based on Table I, it can be observed that the proposed student architectures outperform the teacher network and are better than the other methods on the LSUN dataset. The Head III-B student network, which computes an attention map based on the maximum similarity between input and prototype, has the best classification performance among the other student networks. The Head III-C student network, which computes individual attention maps for the input and the prototype, has similar performance to Head III-B but is computationally more expensive to train due to the additional attention map used. It can be observed that ProtoDNN underperforms severely on the LSUN dataset. A manual inspection of the prototypes from ProtoDNN reveals that the prototypes do not resemble any training samples and that the autoencoder used for prototype decoding fails to reconstruct the LSUN samples, resulting in poor classification performance. This implies that the joint optimization for the layers in ProtoDNN performs well only on simple grayscale datasets such as MNIST and Fashion MNIST [6] but is not generalizable to complex datasets such as LSUN. Our student network does not rely on autoencoders for prototype visualization, since the prototypes are actual training samples selected from the training set. On the basis of our own observation with outlier detection on simpler datasets and the inability of ProtoDNN to generalize to complex datasets such as LSUN, we refrain from using common and simple datasets such as MNIST and Fashion MNIST in our evaluation to provide a more challenging performance benchmark.

For the PCam and Stanford Cars datasets, as shown in Tables II and III, we show only one type of head per category. For the Head II and III categories, we show only performances of Head II-B and Head III-B architectures as they are the general cases of Head II-A and Head III-A, respectively, and have lower computational complexity than the Head III-C

TABLE I
CLASSIFICATION PERFORMANCE ON LSUN TEST SET.

Model	Acc. (%)
χ^2 SVM [42]	76.8
ProtoDNN [6]	54.2
ProtoPNet [8]	82.5
Teacher (<i>baseline</i>)	80.1
Student Head I	82.9
Student Head II-A	82.9
Student Head II-B	82.5
Student Head III-A	82.8
Student Head III-B	83.5
Student Head III-C	83.4

TABLE II
CLASSIFICATION PERFORMANCE ON PCAM TEST SET.

Model	Acc. (%)	AUROC
χ^2 SVM [42]	80.2	0.8875
ProtoDNN [6]	76.7	0.8180
ProtoPNet [8]	83.4	0.8822
Teacher (<i>baseline</i>)	80.5	0.9142
Student Head I	83.3	0.9216
Student Head II-B	82.5	0.9111
Student Head III-B	81.0	0.9205

architecture. Based on Table II, ProtoPNet and our student network with Head I architecture demonstrate equally good performance on the PCam dataset. The former achieves the highest test accuracy (only marginally better than Head I), and the latter achieves the best AUROC score. The PCam dataset consisting of only 2 classes and similar image statistics has smaller intraclass variations than the LSUN dataset with its 10 classes. In this case, Head I, which computes similarity using spatially pooled features, overfits less, thus performing better than more complex architectures such as Head II-B and Head III-B.

In Table III, we compare the different head architectures only with the most competent state-of-the-art model, i.e., ProtoPNet. For the student networks trained on the Stanford Cars dataset, we observe similar performance trend as the student networks trained on PCam dataset since the Stanford Cars classification task is a fine-grained car model classification with significantly smaller intra- and interclass variations (similar to PCam samples) than the LSUN dataset. The architecture Head I has the best classification performance, surpassing the performance of ProtoPNet, while the attention-guided Head III-B architecture with higher complexity underperforms on this dataset. We used smaller learning rates for the Head III-B network on this dataset as the model overfits severely with higher learning rates as observed in our initial experiment. In an attempt to improve the performance of the model on this dataset, we also investigated several Head III-B variants such as a parameterized spatial attention and uniform attention variants but obtain unsatisfactory classification performance compared to the Head I and Head II-B networks. For more information, we refer the reader to Section S-V of the Supplementary Material. Due to small sample variations in the fine-grained dataset and a larger number of learnable parameters in the Head III-B architecture (the additional parameters in the 1D-CNN layer), the model has a higher tendency to overfit

³https://github.com/pennychong94/Project_Towards-Scalable-and-Unified-Example-based-Explanation-and-OD

compared to simpler architectures such as Head I and Head II-B. Since fine-grained classification is a task that involves distinguishing between very similar objects, the network may need to focus on a smaller region of the image sample. Thus, Head III-B can potentially be improved further by sampling small prototype patches from the training set, as inspired by the learning of prototype latent patches in ProtoPNet, which permits the attention map in Head III-B to be applied more precisely over a smaller target region. We note that there is a performance gap between the performance of ProtoPNet in Table III and the performance reported in [8] (Table 1 of their Supplementary Material). The performance gap is due to our experimental setup, which is different from [8] as described earlier in Section IV, and the classification problem of the first 50 classes that are naturally more challenging than the remaining Stanford Cars classes. For a complete discussion on this matter, we refer the reader to Section S-IV of the Supplementary Material.

For the remaining experiments using the PCam and Stanford Cars datasets, we evaluated only Head II-B and Head III-B for the Head II and Head III categories.

TABLE III

CLASSIFICATION PERFORMANCE ON STANFORD CARS TEST SET USING ONLY SAMPLES FROM THE FIRST 50 CLASSES. THE HEAD III-B NETWORK FOR THIS DATASET IS TRAINED USING LOWER LEARNING RATES THAN THE OTHER HEAD ARCHITECTURES. REFER TO SECTION IV FOR MORE DETAILS ON THE EXPERIMENTAL SETTING.

Model	Acc. (%)
ProtoPNet [8]	69.2
Teacher (<i>baseline</i>)	70.0
Student Head I	72.1
Student Head II-B	70.7
Student Head III-B	60.6

TABLE IV

PERFORMANCE OF THE STATE-OF-THE-ART MODELS COMPARED TO OUR TEACHER MODEL ON LSUN AND PCAM TEST SETS. *OUR TEACHER NETWORK IS TRAINED ON A SUBSET OF THE FULL LSUN TRAINING SET USING ONLY 8000 SAMPLES PER CLASS. THE REPORTED RESULTS FOR LSUN ARE THE PERFORMANCES ON THE ORIGINAL VALIDATION SPLIT (CONSISTING OF 300 IMAGES PER CATEGORY) THAT WE USED AS OUR TEST SET.

LSUN Dataset	Acc. (%)
*Teacher (<i>baseline</i>)	80.1
Context-CNN [43]	89.0
SIAT_MMLAB [44]	91.8
PCam Dataset	AUROC
Teacher (<i>baseline</i>)	0.914
VF-CNN [45]	0.951
G-CNN [46], [47]	0.968
DSF-CNN [48]	0.975

We conclude that using prototype-based students generally results in models that are equally competitive in prediction performance as the teacher network except for Head III-B on the Stanford Cars dataset. Employing the same underlying architecture and using soft labels from the teacher network to train the prototype networks allows unconstrained modeling of the original task. We note that our proposed models are built on top of a typical CNN model and thus have lower performance than more advance state-of-the-art methods that are curated to

perform well on the task. To achieve performance comparable to Table IV, one may use any state-of-the-art model as the teacher network and use the soft labels and the same underlying architecture to train the prototype networks. We do not include the state-of-the-art performance on the Stanford Cars dataset in Table IV since we used only the first 50 car classes in our experiments due to time constraints.

B. Top-k Nearest Example and Pixelwise Explanation

We provide an explanation using the top- k nearest prototypical examples from different classes and show the pixel evidence of similarity between the prototype and test sample. We first compare the quality of the pixelwise LRP heatmap explanation generated by the different student architectures with the teacher network and then show examples of explanation in the two forms (refer to Section III-D). We assess

Fig. 5. LRP heatmap quality evaluation on LSUN (left) and PCam (right) datasets. The LRP heatmap perturbation for the student will use the LRP score computed between each test sample and its top-1 prototype.

the quality of the pixelwise LRP heatmap explanation for the network by performing a series of region perturbations in the input space of the test sample. The LRP heatmap comprises relevance scores associated with the degree of importance for the predicted class. Performing a perturbation on the image pixel associated with high relevance scores will destroy the network logit for the predicted class. In other words, we expect a sharp decrease in the logit score of the predicted class as the relevant features in the input sample are gradually removed in decreasing order of importance. We adopt the generalized region perturbation approach in [49]. We perform a series of perturbations starting from the most relevant region (based on the LRP relevance scores) to the least relevant region, as shown in Figure 5. A steep decrease suggests a good pixel explanation.

Based on the subplot for the LSUN dataset in Figure 5, the Head III-B student architecture with the best classification performance also has the best heatmap explanation, as shown by the largest decline in the logit scores. The attention modules in Head III-A and Head III-B show evidently better explanation qualities than the standard teacher network, as suggested by their respective curves. The Head I architecture shows a poorer explanation than the teacher network due to the early average pooling operation, which may have removed distinctive features that are useful in providing an informative pixel explanation. For the PCam dataset, the subplot suggests

Fig. 6. LRP heatmaps from Head III-B on the LSUN dataset for the prediction *classroom* (left) and *dining room* (right). Subplots in the upper part show heatmaps explaining the input sample, while those in the lower part show heatmaps explaining the shown prototype. Each column corresponds to the heatmaps of an input-prototype pair with respect to the top-1 prototype of a specific class. The heatmaps in the second row of the upper and lower parts are unscaled, whereas the heatmaps in the third row are scaled such that the (x, p_k) pair with lower u_k score (written as sim score) will appear dimmer.

Fig. 7. LRP heatmaps from Head III-B on the PCam dataset. The first two columns are for a true negative (tumor absent), the third and fourth columns are for a true positive (tumor present), and the last two columns are for a false negative.

better heatmap quality for Head III-B architecture in the first 5 perturbation steps only and is overtaken by the teacher network and the Head II-B student network. Due to a lower diversity of objects in a patch (small cell types, stromal tissue and background) in the PCam dataset compared to the LSUN dataset, there may be a lesser need for attention weighting for explanation; thus, Head II-B shows better LRP explanation than its attention-augmented counterpart Head III-B when an increasing number of relevant regions are perturbed. Likewise, Head I architecture shows poor LRP explanation on the PCam dataset although it achieves the best classification performance in Table II. Generally, architectures with an attention mechanism provide better LRP explanation for datasets with

large object diversity, while datasets with low object diversity need lesser attention weighting for explanation. In summary, Head II and Head III architectures produce heatmap quality comparable to the teacher network.

We show only LRP heatmaps from architecture with the best explanation (Head III-B) in Figures 6 and 7. We show pairs of heatmaps between the input sample and the top-1 nearest prototype (based on u_k score) for a few selected classes. The positive and negative evidence in the heatmaps are represented by red and blue pixels, respectively. We obtain nonidentical input heatmaps (the columns in the top half of Figures 6 and 7) when comparing the same input sample with different prototype samples due to the different similarity outputs computed and thus unequal relevances for each input-prototype pair. In the left part of Figure 6, heatmaps for mostly all classes are dimmed out (the third row in the upper and lower part) after scaling due to their low similarity u_k scores, except class 3 (*classroom*), which is the predicted class. The red pixels are the similar regions between the pair of input-prototype that has contributed positively to the prediction *classroom*. The input heatmaps in the second row in columns 1, 4, and 5 look similar and they correspond to the similarity of the input image (*classroom*) to prototype examples belonging to the categories bedroom, classroom, and conference room, respectively. The categories bedroom, classroom, and conference room are all indoor scenes containing somewhat similar objects and indoor lighting. Since for each class we look at the nearest prototype to an image (rather than a randomly drawn prototype), we will likely find similar sceneries for these three indoor categories in a large dataset such as LSUN. Thus, the input heatmaps that are computed based on looking at the nearest prototype for each category, appear to reflect a plausible similarity among these three similar indoor scene examples.

In the lower right of Figure 6 showing the top-1 prototypes

from different classes, the regions that resemble furniture in the dining room, such as chairs, tables, and stools, generally lit up in the prototype heatmaps for the prediction *dining room*. They mainly correlate to the input regions with a dining table, chair, and chandelier, as shown in the upper right of Figure 6.

For the explanation on the PCam dataset in Figure 7, we show explanations for a true negative, a true positive, and a false negative. The true negative prediction shows positive relevance scores (in red) on dense clusters of lymphocytes (dark regular nuclei), which is evidence supporting the absence of tumors. For the false negative prediction, the input, which is a positive sample (as indicated by the ground truth), resembles more of the negative class prototype (class 0) rather than the positive class prototype (class 1) in the red pixel areas. One usually looks at the nearest prototype to the test sample for explanation in addition to assessing the heatmap computed for the test sample. In the PCam case, the interpretation of the prototype with its LRP heatmap requires some domain knowledge from the person, such as a pathologist who is trained to recognize certain morphological structures. In Figure

Fig. 8. Input sample and its heatmap for a false negative prediction extracted from the first two rows and the fifth column of Figure 7. The input heatmap is computed based on its similarity with the predicted negative class prototype, as shown in the fifth column of Figure 7.

8, which depicts the false negative explanation extracted from Figure 7 for the input sample with respect to a negative class prototype, one can see in the upper left of the first row (displaying the input sample) a few faint likely cancer nuclei. The heatmap shown in Figure 8 shows very low heatmaps score in these cancer nuclei locations (almost black) and high heatmap scores in their direct surrounding embedding tissue matrix. Moreover, one sees high heatmap scores in some of the black dots, which are tissue invading lymphocytes (TiLs) also present in the prototype, and for which some but not all of them seem to be similar to the TiLs in the prototype (Figure 7, column 5, row 3). By looking at the prototype and its heatmap in the lower part of Figure 7, column 5, one can see that the similarities are found in a part of the tissue invading lymphocytes and a part of the embedding stromal tissue matrix.

The explanations in Figures 6 and 7 allow domain experts to easily validate the reasoning process of the neural networks by verifying that the nearest prototypical training samples are correct and the contributing features are relevant to the predicted class. By analyzing contradictory cases of prediction

and explanation, domain experts understand the model limitation and are able to make an informed judgment. We have included more LRP explanation examples in Section S-III of the Supplementary Material.

The prototype similarity scores u_k , that are used to identify the top- k nearest prototypical examples can also be employed to gauge the explanation confidence of a given input sample. The optimal value of k that we choose to use for explanation depends on the task at hand. We can consider an explanation based on prototypical examples to be of sufficient confidence if the top- k nearest prototypes (i.e., top- k highest u_k score) from the predicted class are responsible for a high fixed ratio (e.g., 90%) of the prediction mass; thus, these prototypes can be considered to have sufficiently high similarity scores. The top- k nearest prototypes are considered to have insufficient similarity scores if their total score is below this fixed ratio.

Alternatively, one can also interpret the similarity scores in another manner. We consider the set of top- k nearest prototypes from the predicted class to be too small if the removal of the prototype with the smallest similarity score would change the prediction label such that it differed from the prediction with the full set of prototypes. This is related to the principle of pertinent positives [50], but performed on the prototype level rather than the level of subsets of an input sample.

Fig. 9. Impact on performance using different numbers of prototypes. The x-axis is the total number of prototypes from all classes.

C. Detecting Outliers for Additional Validation

The prototype student networks can naturally quantify the degree of outlieriness using the set of similarity scores. The GODIN method shows the best outlier detection performance in the flowers setup, as shown in Table V, due to significant differences in the data distribution between the LSUN and flowers datasets. GODIN does not perform well on more challenging outlier setups such as strokes and altered color setups in which anomalies are created via the manipulation of normal samples.

The Head I student architecture in Table V achieves the second-best performance for the flowers setup and top performance for the altered color setup using the top-20 prototype similarity scores u_k . Since the image statistics of the flowers dataset and the altered color anomalies are substantially different from the normal LSUN samples, the use of early spatial pooling before computing similarity in Head I does not

TABLE V

OUTLIER DETECTION PERFORMANCE REPORTED USING THE **AUROC|FPR95** METRIC FOR LSUN NETWORKS. HIGHER ($\uparrow$) AUROC VALUES INDICATE BETTER OUTLIER DETECTION PERFORMANCE, WHEREAS LOWER ($\downarrow$) FPR95 VALUES INDICATE BETTER OUTLIER DETECTION PERFORMANCE.

Model	Setup A Flowers			Setup B LSUN strokes_5			Setup C LSUN altered color		
	<i>Anomaly Score</i> *			<i>Anomaly Score</i> *			<i>Anomaly Score</i> *		
IF [39]	0.678 0.661			0.687 0.811			0.482 0.936		
Teacher ¹ + GODIN [27]	0.992 0.041			0.715 0.759			0.695 0.807		
Teacher ² + GODIN [27]	0.995 0.019			0.690 0.786			0.692 0.824		
	<i>Max Prob.</i> [25]			<i>Max Prob.</i> [25]			<i>Max Prob.</i> [25]		
χ^2 SVM [42]	0.921 0.313			0.742 0.737			0.691 0.820		
ProtoDNN [6]	0.553 0.935			0.497 0.961			0.520 0.913		
ProtoPNet [8]	0.847 0.478			0.645 0.898			0.683 0.847		
Teacher (<i>baseline</i>)	0.934 0.284			0.677 0.800			0.690 0.798		
Teacher + OE [28] (flowers)	NA			0.679 0.812			0.653 0.868		
Teacher + OE [28] (strokes_5)	0.911 0.274			NA			0.710 0.821		
Teacher + OE [28] (color)	0.949 0.179			0.663 0.890			NA		
	<i>Top-1</i>	<i>Top-20</i>	<i>All proto.</i>	<i>Top-1</i>	<i>Top-20</i>	<i>All proto.</i>	<i>Top-1</i>	<i>Top-20</i>	<i>All proto.</i>
Student Head I	0.990 0.042	0.993 0.025	0.739 0.511	0.712 0.832	0.692 0.846	0.473 0.928	0.782 0.774	0.852 0.593	0.630 0.734
Student Head II-A	0.970 0.129	0.965 0.151	0.786 0.632	0.764 0.719	0.738 0.799	0.611 0.858	0.809 0.661	0.835 0.631	0.710 0.826
Student Head II-B	0.972 0.151	0.984 0.082	0.986 0.068	0.763 0.749	0.765 0.793	0.677 0.902	0.779 0.770	0.790 0.799	0.778 0.892
Student Head III-A	0.922 0.315	0.943 0.221	0.885 0.328	0.726 0.748	0.735 0.701	0.683 0.772	0.672 0.822	0.691 0.766	0.636 0.853
Student Head III-B	0.855 0.657	0.906 0.542	0.909 0.434	0.692 0.837	0.702 0.842	0.718 0.766	0.674 0.852	0.707 0.834	0.731 0.772
Student Head III-C	0.970 0.118	0.972 0.071	0.914 0.434	0.726 0.769	0.727 0.772	0.509 0.890	0.763 0.763	0.786 0.690	0.680 0.804

* The anomaly detection algorithm returns an anomaly score for every input sample, and such an algorithm is not used for classification.

¹ The parameters in the trained layers of the teacher network are fixed, and optimization involves only the GODIN layers, i.e., $g(x)$ and $h_i(x)$.

² The parameters in the trained layers of the teacher network are fine-tuned with a learning rate that is an order smaller than the learning rate for the GODIN layers, i.e., $g(x)$ and $h_i(x)$.

TABLE VI

OUTLIER DETECTION PERFORMANCE REPORTED USING THE **AUROC|FPR95** METRIC FOR PCAM NETWORKS. HIGHER ($\uparrow$) AUROC VALUES INDICATE BETTER OUTLIER DETECTION PERFORMANCE, WHEREAS LOWER ($\downarrow$) FPR95 VALUES INDICATE BETTER OUTLIER DETECTION PERFORMANCE.

Model	Setup B PCam strokes_5			Setup C PCam altered color		
	<i>Anomaly Score</i> *			<i>Anomaly Score</i> *		
IF [39]	0.723 0.712			0.580 0.754		
Teacher ¹ + GODIN [27]	0.696 0.824			0.585 0.967		
Teacher ² + GODIN [27]	0.579 0.929			0.788 0.857		
	<i>Max Prob.</i> [25]			<i>Max Prob.</i> [25]		
χ^2 SVM [42]	0.521 0.938			0.519 0.933		
ProtoDNN [6]	0.441 0.976			0.087 0.991		
ProtoPNet [8]	0.514 0.930			0.450 0.654		
Teacher (<i>baseline</i>)	0.527 0.905			0.326 0.910		
Teacher + OE [28] (strokes_5)	NA			0.703 0.640		
Teacher + OE [28] (color)	0.561 1.00			NA		
	<i>Top-1</i>	<i>Top-20</i>	<i>All proto.</i>	<i>Top-1</i>	<i>Top-20</i>	<i>All proto.</i>
Student Head I	0.576 0.817	0.574 0.816	0.511 0.960	0.350 0.822	0.358 0.853	0.820 0.427
Student Head II-B	0.678 0.753	0.661 0.730	0.636 0.882	0.377 0.929	0.394 0.870	0.823 0.325
Student Head III-B	0.687 0.778	0.631 0.829	0.541 0.835	0.329 0.908	0.249 0.955	0.224 0.959

* The anomaly detection algorithm returns an anomaly score for every input sample, and such an algorithm is not used for classification.

¹ The parameters in the trained layers of the teacher network are fixed, and optimization involves only the GODIN layers, i.e., $g(x)$ and $h_i(x)$.

² The parameters in the trained layers of the teacher network are fine-tuned with a learning rate that is an order smaller than the learning rate for the GODIN layers, i.e., $g(x)$ and $h_i(x)$.

remove distinctive features of the anomalies but is effective in detecting outliers while offering the least potential for overfitting. For the strokes outlier setup, Head II-B generally has the best detection performance, as shown in Table V and Table I of the Supplementary Material. The placement of the spatial pooling operation before the similarity operation in Head I may have removed traces of strokes on the image, while Head II-B retains the traces of strokes during the sim-

ilarity computation. The Head III-B architecture has a higher architectural complexity than Head I and Head II and is thus less robust to outliers and prone to overfitting. We show the outlier performance of LSUN networks under the AUPRout and AUPRin metrics in Section S-II of the Supplementary Material due to space constraints.

We report detection only on the more challenging setups (strokes and altered color as suggested by the average detection

TABLE VII

OUTLIER DETECTION PERFORMANCE REPORTED USING THE AUROC|FPR95 METRIC FOR STANFORD CARS NETWORKS. HIGHER ($\uparrow$) AUROC VALUES INDICATE BETTER OUTLIER DETECTION PERFORMANCE, WHEREAS LOWER ($\downarrow$) FPR95 VALUES INDICATE BETTER OUTLIER DETECTION PERFORMANCE.

Model	Setup B Stanford Cars strokes_5			Setup C Stanford Cars altered color		
IF [39] Teacher ¹ + GODIN [27] Teacher ² + GODIN [27]	<i>Anomaly Score</i> *			<i>Anomaly Score</i> *		
	0.874 0.533			0.854 0.531		
	0.622 0.940			0.503 0.994		
ProtoPNet [8] Teacher (<i>baseline</i>) Teacher + OE [28] (strokes_5) Teacher + OE [28] (color)	<i>Max Prob.</i> [25]			<i>Max Prob.</i> [25]		
	0.728 1.000			0.792 0.592		
	0.760 0.725			0.789 0.635		
Student Head I Student Head II-B Student Head III-B	NA			0.998 0.100		
	0.851 0.569			NA		
	<i>Top-1</i>	<i>Top-20</i>	<i>All proto.</i>	<i>Top-1</i>	<i>Top-20</i>	<i>All proto.</i>
Student Head I	0.815 0.706	0.808 0.723	0.627 0.844	0.838 0.675	0.859 0.551	0.725 0.729
Student Head II-B	0.818 0.687	0.799 0.770	0.393 0.953	0.810 0.662	0.782 0.741	0.265 0.975
Student Head III-B	0.700 0.857	0.661 0.899	0.551 0.890	0.913 0.361	0.931 0.307	0.903 0.346

* The anomaly detection algorithm returns an anomaly score for every input sample, and such an algorithm is not used for classification.

¹ The parameters in the trained layers of the teacher network are fixed, and optimization involves only the GODIN layers, i.e., $g(x)$ and $h_i(x)$.

² The parameters in the trained layers of the teacher network are fine-tuned with a learning rate that is an order smaller than the learning rate for the GODIN layers, i.e., $g(x)$ and $h_i(x)$.

TABLE VIII

CLASSIFICATION PERFORMANCE WITHOUT AND WITH PRUNING.

Model	Prune	LSUN	PCam	
		Acc. (%)	Acc. (%)	AUROC
Head I	$\times$	82.9	83.3	0.9216
	$\checkmark$	82.7	82.6	0.9177
Head II-B	$\times$	82.5	82.5	0.9111
	$\checkmark$	82.6	83.3	0.9111
Head III-B	$\times$	83.5	81.0	0.9205
	$\checkmark$	83.2	81.8	0.9244

scores in Table V) for the PCam and Stanford Cars datasets.

In Table VI, the IF method generally achieves the best detection results on the strokes setup. Likewise, the student architectures with similarity layers before averaging (Head II-B and Head III-B) perform better in detecting strokes artifacts on the PCam dataset, which is consistent with the LSUN networks. For the altered color setup, as shown in Table VI, student networks with lower architectural complexity, such as Head I and Head II-B architectures, perform well on the altered color setup compared to Head III-B with higher architectural complexity. The Head II-B and Head I networks using all prototypes are equally competent on the PCam altered color setting. The outliers created using the PCam dataset in this setting are easier to be detected than those created using the LSUN dataset due to the significant differences between the color distribution of outliers and normal PCam samples. We also observe that the ProtoDNN underperforms on the altered color setup, as a majority of the outliers have a higher maximum class probability than the normal samples. We show the outlier performance of PCam networks under the AUPRout and AUPRin metrics in Section S-II of the Supplementary Material due to space constraints.

For the Stanford Cars experiments in Table VII, we also observe that the IF model performs best in the strokes setup, similar to the PCam experiments. The competing OE method

achieves comparable performance to IF on the Stanford Cars strokes setup and superior performance on the altered color setup. We deduce that exposing the network to an outlier set during training when the in-distribution set has low variability among samples can be highly beneficial to the model, as it is able to learn a more conservative concept of inliers easily as opposed to a setting with highly varied in-distribution samples. We also observe that the Head III-B network with poor classification performance in Table III performs better than the other head architectures in the altered color setup but not in the strokes setup as opposed to its performance on PCam outliers in Table VI. Due to lack of space, we report the outlier performance of Stanford Cars networks under the AUPRout and AUPRin metrics in Section S-II of the Supplementary Material.

In conclusion, except for the PCam altered color setup, the use of 20 prototypes for Head II and Head III architectures generally shows reasonable outlier performances, and are mostly better than or comparable to the IF, GODIN, and OE methods. This supports our argument for unified prototype-based models with generic head architectures that are interpretable and robust to outliers.

VI. DISCUSSION

A. Varying Number of Prototypes

We investigate the impact on classification performance using different numbers of prototypes. We use an equal number of prototypes per class. We discuss only the performance for student networks with simple (Head I) and complex (Head III-B) architectures due to a lack of space. Based on Figure 9, only a marginal change in performance can be seen for the simple Head I architecture on the LSUN dataset, whereas a huge decline in performance can be seen for the more complex Head III-B architecture when using a large number of prototypes (150 prototypes) due to its larger tendency to overfit. Likewise,

TABLE IX
OUTLIER DETECTION PERFORMANCE REPORTED USING THE AUROC METRIC FOR NETWORKS WITHOUT AND WITH PRUNING.

Model	Prune	LSUN									PCam					
		Setup A Flowers			Setup B LSUN strokes_5			Setup C LSUN altered color			Setup B PCam strokes_5			Setup C PCam altered color		
		<i>Top-1</i>	<i>Top-20</i>	<i>All proto.</i>	<i>Top-1</i>	<i>Top-20</i>	<i>All proto.</i>	<i>Top-1</i>	<i>Top-20</i>	<i>All proto.</i>	<i>Top-1</i>	<i>Top-20</i>	<i>All proto.</i>	<i>Top-1</i>	<i>Top-20</i>	<i>All proto.</i>
Head I	✗	0.990	0.993	0.739	0.712	0.692	0.473	0.782	0.852	0.630	0.576	0.574	0.511	0.350	0.358	0.820
	✓	0.986	0.982	0.704	0.701	0.626	0.477	0.773	0.821	0.635	0.579	0.506	0.467	0.367	0.323	0.584
Head II-B	✗	0.972	0.984	0.986	0.763	0.765	0.677	0.779	0.790	0.778	0.678	0.661	0.636	0.377	0.394	0.823
	✓	0.968	0.983	0.977	0.749	0.765	0.672	0.757	0.812	0.777	0.673	0.701	0.676	0.443	0.552	0.825
Head III-B	✗	0.855	0.906	0.909	0.692	0.702	0.718	0.674	0.707	0.731	0.687	0.631	0.541	0.329	0.249	0.224
	✓	0.821	0.948	0.933	0.689	0.754	0.742	0.677	0.780	0.774	0.641	0.622	0.623	0.234	0.211	0.257

only slight performance differences are reported for the PCam dataset with different numbers of prototypes. For the relatively simple PCam dataset, using a small number of prototypes (20 prototypes) is sufficient and less likely to overfit, as shown by the high test accuracy for both architectures. Hence, using a small number of prototypes for simple datasets and complex network architectures is a good practice.

B. Pruning of Prototypes and Relevant Neurons

Since the prototype importance weight $\mathcal{M}$ defines the importance of the K prototypes in the network prediction, we can prune the set of prototypes with the least weight in $\mathcal{M}$ to evaluate the impact on classification and outlier performances. We do not evaluate the LRP explanation of the pruned networks, as the general architecture of the networks remains largely unchanged. We prune $p = 30\%$ (the same as p used in the iterative prototype replacement algorithm) of the prototypes and the relevant neurons, and then we fine-tune the network. Alternatively, one can set the pruning factor based on the classification performance on the validation set.

Based on Table VIII, we observe marginal changes in the classification performance after pruning. The pruned network with Head II-B shows a slight increase in accuracy for both datasets. The outlier detection performance of the pruned networks is reported in Table IX. The pruned networks exhibit a larger change in the outlier performance compared to the prediction accuracy. For the LSUN dataset, the student network with the Head III-B architecture consistently performs better with pruning, as indicated by the significant gain in the detection performance. Since the pruning of prototypes and relevant neurons reduces the size of the network, the Head III-B network overfits less after fine-tuning and is more robust to outliers. We do not observe the same benefit of pruning for Head III-B on the PCam outlier setups. For the strokes setup, an evident gain in performance can be observed for the Head II-B network trained on the PCam dataset. A smaller gain is seen for the Head II-B network in the altered color setup for both datasets. Thus, one may consider pruning uninformative prototypes if the network is prone to overfitting.

VII. CONCLUSION

Various student architectures that compute prototype-based predictions were investigated for their suitability to provide example-based explanations and detect outliers. The selection

of prototypical examples from the training set via the iterative prototype replacement method guaranteed meaningful explanation. The prototype similarity scores can also be used naturally to quantify the outlierness of a sample. However, the architecture with the best LRP explanation did not achieve the best outlier detection performance.

REFERENCES

- [1] K. Simonyan, A. Vedaldi, and A. Zisserman, “Deep inside convolutional networks: Visualising image classification models and saliency maps,” *arXiv preprint arXiv:1312.6034*, 2013.
- [2] M. D. Zeiler and R. Fergus, “Visualizing and understanding convolutional networks,” in *European conference on computer vision*. Springer, 2014, pp. 818–833.
- [3] G. Montavon, A. Binder, S. Lapuschkin, W. Samek, and K.-R. Müller, “Layer-wise relevance propagation: an overview,” in *Explainable AI: interpreting, explaining and visualizing deep learning*. Springer, 2019, pp. 193–209.
- [4] C. Cortes and V. Vapnik, “Support-vector networks,” *Machine learning*, vol. 20, no. 3, pp. 273–297, 1995.
- [5] H.-M. Yang, X.-Y. Zhang, F. Yin, and C.-L. Liu, “Robust classification with convolutional prototype learning,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 3474–3482.
- [6] O. Li, H. Liu, C. Chen, and C. Rudin, “Deep learning for case-based reasoning through prototypes: A neural network that explains its predictions,” *arXiv preprint arXiv:1710.04806*, 2017.
- [7] Y. Ming, P. Xu, H. Qu, and L. Ren, “Interpretable and steerable sequence learning via prototypes,” in *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2019, pp. 903–913.
- [8] C. Chen, O. Li, D. Tao, A. Barnett, C. Rudin, and J. K. Su, “This looks like that: deep learning for interpretable image recognition,” in *NeurIPS*, 2019, pp. 8930–8941.
- [9] P. Hase, C. Chen, O. Li, and C. Rudin, “Interpretable image recognition with hierarchical prototypes,” in *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, vol. 7, no. 1, 2019, pp. 32–40.
- [10] D. Rymarczyk, Ł. Struski, J. Tabor, and B. Zieliński, “Protopshare: Prototype sharing for interpretable image classification and similarity discovery,” *arXiv preprint arXiv:2011.14340*, 2020.
- [11] M. Nauta, R. van Bree, and C. Seifert, “Neural prototype trees for interpretable fine-grained image recognition,” *arXiv preprint arXiv:2012.02046*, 2020.
- [12] K. S. Gurumoorthy, A. Dhurandhar, G. Cecchi, and C. Aggarwal, “Efficient data representation by selecting prototypes with importance weights,” in *2019 IEEE International Conference on Data Mining (ICDM)*. IEEE, 2019, pp. 260–269.
- [13] B. Kim, O. Koyejo, R. Khanna *et al.*, “Examples are not enough, learn to criticize! criticism for interpretability,” in *NIPS*, 2016, pp. 2280–2288.
- [14] S. O. Arık and T. Pfister, “Protoattend: Attention-based prototypical learning,” *Journal of Machine Learning Research*, vol. 21, pp. 1–35, 2020.
- [15] N. S. Altman, “An introduction to kernel and nearest-neighbor non-parametric regression,” *The American Statistician*, vol. 46, no. 3, pp. 175–185, 1992.

- [16] A. Chattopadhyay, A. Sarkar, P. Howlader, and V. N. Balasubramanian, "Grad-cam++: Generalized gradient-based visual explanations for deep convolutional networks," in *2018 IEEE winter conference on applications of computer vision (WACV)*. IEEE, 2018, pp. 839–847.
- [17] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," in *Proceedings of the 31st international conference on neural information processing systems*, 2017, pp. 4768–4777.
- [18] J. V. Jeyakumar, J. Noor, Y.-H. Cheng, L. Garcia, and M. Srivastava, "How can i explain this to you? an empirical study of deep neural network explanation methods," *Advances in Neural Information Processing Systems*, 2020.
- [19] F. Angiulli and C. Pizzuti, "Fast outlier detection in high dimensional spaces," in *Principles of Data Mining and Knowledge Discovery*, T. Elomaa, H. Mannila, and H. Toivonen, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2002, pp. 15–27.
- [20] M. M. Breunig, H.-P. Kriegel, R. T. Ng, and J. Sander, "Lof: Identifying density-based local outliers," in *ACM SIGMOD*, 2000, p. 93–104.
- [21] L. J. Latecki, A. Lazarevic, and D. Pokrajac, "Outlier detection with kernel density functions," in *MLDM 2007*, ser. Lecture Notes in Computer Science, P. Perner, Ed., vol. 4571. Springer, 2007, pp. 61–75.
- [22] S. Kim and S. Cho, "Prototype based outlier detection," in *The 2006 IEEE International Joint Conference on Neural Network Proceedings*. IEEE, 2006, pp. 820–826.
- [23] R. Ramirez-Padron, D. Foregger, J. Manuel, M. Georgiopoulos, and B. Mederos, "Similarity kernels for nearest neighbor-based outlier detection," in *International Conference on Advances in Intelligent Data Analysis*. Berlin, Heidelberg: Springer-Verlag, 2010, p. 159–170.
- [24] X. Gu, L. Akoglu, and A. Rinaldo, "Statistical analysis of nearest neighbor methods for anomaly detection," in *NeurIPS*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, Eds., 2019, pp. 10 923–10 933. [Online]. Available: <http://papers.nips.cc/paper/9274-statistical-analysis-of-nearest-neighbor-methods-for-anomaly-detection.pdf>
- [25] D. Hendrycks and K. Gimpel, "A baseline for detecting misclassified and out-of-distribution examples in neural networks," *arXiv preprint arXiv:1610.02136*, 2016.
- [26] S. Liang, Y. Li, and R. Srikant, "Enhancing the reliability of out-of-distribution image detection in neural networks," *arXiv preprint arXiv:1706.02690*, 2017.
- [27] Y.-C. Hsu, Y. Shen, H. Jin, and Z. Kira, "Generalized odin: Detecting out-of-distribution image without learning from out-of-distribution data," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 10 951–10 960.
- [28] D. Hendrycks, M. Mazeika, and T. Dietterich, "Deep anomaly detection with outlier exposure," *arXiv preprint arXiv:1812.04606*, 2018.
- [29] G. Hinton, O. Vinyals, and J. Dean, "Distilling the knowledge in a neural network," *arXiv preprint arXiv:1503.02531*, 2015.
- [30] D. Gong, L. Liu, V. Le, B. Saha, M. R. Mansour, S. Venkatesh, and A. v. d. Hengel, "Memorizing normality to detect anomaly: Memory-augmented deep autoencoder for unsupervised anomaly detection," in *Proceedings of the IEEE International Conference on Computer Vision*, 2019, pp. 1705–1714.
- [31] S. Bach, A. Binder, G. Montavon, F. Klauschen, K.-R. Müller, and W. Samek, "On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation," *PloS one*, vol. 10, no. 7, p. e0130140, 2015.
- [32] O. Eberle, J. Büttner, F. Kräutli, K.-R. Müller, M. Valleriani, and G. Montavon, "Building and interpreting deep similarity models," *arXiv preprint arXiv:2003.05431*, 2020.
- [33] F. Yu, A. Seff, Y. Zhang, S. Song, T. Funkhouser, and J. Xiao, "Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop," *arXiv preprint arXiv:1506.03365*, 2015.
- [34] B. E. Bejnordi, M. Veta, P. J. Van Diest, B. Van Ginneken, N. Karssemeijer, G. Litjens, J. A. Van Der Laak, M. Hermans, Q. F. Manson, M. Balkenhol *et al.*, "Diagnostic assessment of deep learning algorithms for detection of lymph node metastases in women with breast cancer," *Jama*, vol. 318, no. 22, pp. 2199–2210, 2017.
- [35] B. S. Veeling, J. Linmans, J. Winkens, T. Cohen, and M. Welling, "Rotation equivariant cnns for digital pathology," in *MICCAI*. Springer, 2018, pp. 210–218.
- [36] J. Krause, M. Stark, J. Deng, and L. Fei-Fei, "3d object representations for fine-grained categorization," in *4th International IEEE Workshop on 3D Representation and Recognition (3dRR-13)*, Sydney, Australia, 2013.
- [37] M.-E. Nilsback and A. Zisserman, "Automated flower classification over a large number of classes," in *2008 Sixth Indian Conference on Computer Vision, Graphics & Image Processing*. IEEE, 2008, pp. 722–729.
- [38] P. Chong, L. Ruff, M. Kloft, and A. Binder, "Simple and effective prevention of mode collapse in deep one-class classification," in *IJCNN 2020*, 2020, pp. 1–9.
- [39] F. Liu, K. Ting, and Z. Zhou, "Isolation forest," in *2008 Eighth IEEE International Conference on Data Mining*. IEEE, 2008, pp. 413–422.
- [40] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [41] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer, "Automatic differentiation in pytorch," 2017.
- [42] A. Vedaldi and A. Zisserman, "Efficient additive kernels via explicit feature maps," *IEEE transactions on pattern analysis and machine intelligence*, vol. 34, no. 3, pp. 480–492, 2012.
- [43] S. A. Javed and A. K. Nelakanti, "Object-level context modeling for scene classification with context-cnn," *arXiv preprint arXiv:1705.04358*, 2017.
- [44] L. Wang, S. Guo, W. Huang, Y. Xiong, and Y. Qiao, "Knowledge guided disambiguation for large-scale scene classification with multi-resolution cnns," *IEEE Transactions on Image Processing*, vol. 26, no. 4, pp. 2055–2068, 2017.
- [45] D. Marcos, M. Volpi, N. Komodakis, and D. Tuia, "Rotation equivariant vector field networks," in *Proceedings of the IEEE International Conference on Computer Vision*, 2017, pp. 5048–5057.
- [46] E. J. Bekkers, M. W. Lafarge, M. Veta, K. A. Eppenhof, J. P. Pluim, and R. Duits, "Roto-translation covariant convolutional networks for medical image analysis," in *International conference on medical image computing and computer-assisted intervention*. Springer, 2018, pp. 440–448.
- [47] M. W. Lafarge, E. J. Bekkers, J. P. Pluim, R. Duits, and M. Veta, "Roto-translation equivariant convolutional networks: Application to histopathology image analysis," *Medical Image Analysis*, vol. 68, p. 101849, 2021.
- [48] S. Graham, D. Epstein, and N. Rajpoot, "Dense steerable filter cnns for exploiting rotational symmetry in histology images," *IEEE Transactions on Medical Imaging*, vol. 39, no. 12, pp. 4124–4136, 2020.
- [49] W. Samek, A. Binder, G. Montavon, S. Lapuschkin, and K.-R. Müller, "Evaluating the visualization of what a deep neural network has learned," *IEEE TNNLS*, vol. 28, no. 11, pp. 2660–2673, 2016.
- [50] A. Dhurandhar, P.-Y. Chen, R. Luss, C.-C. Tu, P. Ting, K. Shanmugam, and P. Das, "Explanations based on the missing: Towards contrastive explanations with pertinent negatives," *arXiv preprint arXiv:1802.07623*, 2018.

Penny Chong is a research scientist at IBM Research, Singapore. She received her Ph.D. in Information Systems Technology and Design from Singapore University of Technology and Design (SUTD) in 2021, under the supervision of Alexander Binder and Ngai-Man Cheung. She received a B.Sc. (Hons) in Applied Mathematics with Computing from Universiti Tunku Abdul Rahman (UTAR), Malaysia in 2016. Her research interests include machine learning and explainable AI and its applications.

Ngai-Man Cheung is an associate professor at the Singapore University of Technology and Design (SUTD). He received his Ph.D. degree in Electrical Engineering from the University of Southern California (USC), Los Angeles, CA, in 2008. His research interests include image and signal processing, computer vision and AI.

Yuval Elovici is the director of the Telekom Innovation Laboratories at Ben-Gurion University of the Negev (BGU), head of the BGU Cyber Security Research Center and a professor in the Department of Software and Information Systems Engineering at BGU. He holds a Ph.D. in Information Systems from Tel-Aviv University. His primary research interests are computer and network security, cyber security, web intelligence, information warfare, social network analysis, and machine learning.

Alexander Binder is an associate professor at the University of Oslo (UiO). He received a Dr.rer.nat. in Computer Science from Technische Universität Berlin, Germany in 2013. His research interests include explainable deep learning and medical imaging.