



HAL
open science

Sanity checks and improvements for patch visualisation in prototype-based image classification

Romain Xu-Darme, Georges Quénot, Zakaria Chihani, Marie-Christine
Rousset

► To cite this version:

Romain Xu-Darme, Georges Quénot, Zakaria Chihani, Marie-Christine Rousset. Sanity checks and improvements for patch visualisation in prototype-based image classification. 2023. cea-04253851v1

HAL Id: cea-04253851

<https://cea.hal.science/cea-04253851v1>

Preprint submitted on 17 Jan 2023 (v1), last revised 24 Oct 2023 (v3)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Sanity checks and improvements for patch visualisation in prototype-based image classification

Romain Xu-Darme^{1,2}, Georges Quénot², Zakaria Chihani¹, Marie-Christine Rousset²

¹ Université Paris-Saclay, CEA, List, F-91120, Palaiseau, France

² Univ. Grenoble Alpes, CNRS, Grenoble INP, LIG, F-38000 Grenoble, France

{romain.xu-darme, zakaria.chihani(at)cea.fr}

{georges.quentot, marie-christine.rousset(at)imag.fr}

Abstract

In this work, we perform an in-depth analysis of the visualisation methods implemented in two popular self-explaining models for visual classification based on prototypes - ProtoPNet and ProtoTree. Using two fine-grained datasets (CUB-200-2011 and Stanford Cars), we first show that such methods do not correctly identify the regions of interest inside of the images, and therefore do not reflect the model behaviour. Secondly, using a deletion metric we demonstrate quantitatively that saliency methods such as Smoothgrads or PRP provide more faithful image patches. We also propose a new relevance metric based on the segmentation of the object provided in some datasets (e.g. CUB-200-2011) and show that the imprecise patch visualisations generated by ProtoPNet and ProtoTree can create a false sense of bias that can be mitigated by the use of more faithful methods. Finally, we discuss the implications of our findings for other prototype-based models sharing the same visualisation method.

1. Introduction

During the last decade, the field of Explainable AI (XAI) has progressively gained wide-spread recognition among the scientific community [1, 9, 11, 29]. This evolution reflects the growing need by society for more transparent and accountable systems, especially in critical domains (e.g. autonomous driving, medical diagnosis), as demonstrated by the recent adoption of the European General Data Protection Regulation (GDPR) and the proposal for a European AI Regulation Act.

XAI aims at unravelling the decision-making process of an autonomous system, and encompasses two approaches that are sometimes conflicting [33], yet often complementary. *Post-hoc* explanation methods apply to pre-existing models

(i.e. pre-trained models in the context of machine-learning) and provide various information regarding either a particular decision (*local explanation*) or the general model behaviour (*global explanation*): understandable local approximation [26, 32], saliency maps [10, 38, 39, 41, 43], counterfactuals [15, 24, 25, 45], concept-based explanations [22, 48], etc. In particular, *post-hoc* explanation methods can help explain the decision-making process of Deep Neural Networks (DNNs), which are highly efficient yet highly complex systems. In the field of computer vision for instance, such methods can help identify the attention region of the DNN w.r.t. a particular decision, sometimes highlighting a possible bias in the model [32]. However, it is worth mentioning that this attention region does not - by itself - necessarily explain the model's decision. Indeed, a DNN might be focusing on a relevant part of the image, and yet might produce an incorrect decision, based on high-dimensional features that are inherently abstract due to the network training procedure.

In order to mitigate this limitation in the explainability of complex systems such as DNNs, a second major avenue of research in the field of XAI consists in developing architectures and training procedures such that the resulting model should be *explainable-by-design* (such models are also called *self-explaining* or *interpretable-by-design*). In computer vision, such architectures primarily use either a case-based reasoning mechanism [12, 16, 30, 34, 35] - where new instances of a problem are solved using comparison with examples (*prototypes*) extracted from the training dataset - or concept-based attribution [3, 13]. In particular, ProtoPNet [12] and ProtoTree [30] have shown that explainable-by-design architectures can reach performance levels on par with non-interpretible models on fine-grained recognition tasks [47, 49]. During training, both models extract reference vectors in the latent space of a deep convolutional neural network (DCNN), corresponding to patches of

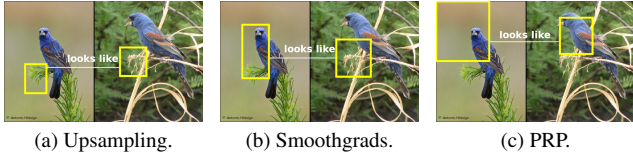


Figure 1. Explanations (“This looks like that”) during inference of a ProtoTree when using upsampling (a), Smoothgrads (b) and PRP (c) for the visualisation of image patches. Due to the imprecision of upsampling when visualising both the prototype (right) and the part in the test image (left), the user might deduce that the model is comparing tree branches, when it is actually also taking the bird into account.

images in the training set (the so-called *prototypes*). During inference, the similarity between a test image and a prototype is evaluated by computing the distance between their respective representations in the latent space of the DCNN. Finally, ProtoPNet and ProtoTree base their decision using the prototypes that show the highest similarity with the test image, and produce explanations by displaying patches of the test image and their most “similar” prototypes side-by-side (*this looks like that* [12]).

However, in practice, both ProtoPNet and ProtoTree sometimes provide explanations using image patches that seem to be focused on the background or elements unrelated to the object itself, indicating a possible bias in the decision (see Fig. 1a). Moreover, while a patch from the test image focusing on the background might invalidate one particular inference, a prototype focusing on background might indicate a more *systemic* bias in the model and seriously hinder its practical accuracy. Although such issue may undermine the trust in self-explaining models, there exists in reality a fundamental imprecision in the patch visualisation methods implemented by ProtoPNet and ProtoTree to generate explanations. In particular, [17] has shown that this imprecision in ProtoPNet visualisation method can sometimes hide a specific type of bias - known as the “Clever Hans” phenomenon - where a given class is highly correlated to a specific visual artefact unrelated to the task at hand (*e.g.* watermarking in an image). More generally, imprecise visualisation methods may suggest model bias where there is none, while sometimes hiding systemic issues of the model.

Our contribution: In this work, we perform an analysis of the visualisation methods implemented in ProtoPNet and ProtoTree, attempting to answer the following research questions: do these architectures generate faithful image patches corresponding to the latent representations used in their decision-making process? do they produce decision based on relevant parts of the image or based on elements of the background? Using two fine-grained datasets (CUB-200-2011 [47] and Stanford Cars [49]) and two saliency

methods (Smoothgrads [41] and PRP [17]), we confirm the results of [17] on ProtoPNet and show that ProtoTree also generates imprecise visual patches. Additionally, using the object segmentation provided in the CUB-200-2011 dataset, we propose a new relevance metric and show quantitatively that in both architectures, such imprecise visualisations often create a false sense of bias that is largely mitigated by the use of more faithful methods. Finally, we discuss the implications of our findings to other prototype-based models sharing the same visualisation method.

This paper is organized as follows: Section 2 describes the related work; Section 3 recalls the theoretical background of ProtoPNet and ProtoTree and introduces the metrics used to evaluate the fidelity and relevance of prototype-based explanation; Section 4 describes the results of our experiments. Finally, Section 5 concludes this contribution and proposes possible improvements to self-explaining models.

2. Related work

Prototype-based classifiers Due to the difficulty of quantifying similarity between images in the visual space (*i.e.* the space of RGB images), self-explaining image classification models based on prototypes first encode images into a high dimensional feature space (also called latent space) - usually using a pre-trained DCNN (*e.g.* Resnet50 [19], VGG [40]) called a backbone, assuming that such encoding preserves visual cues (*e.g.* colour, shape) while being insensitive to minor deformations such as rotations, shifts, changes in scale. This latent representation can be composed of a single vector summarizing the information of the entire image, or composed of an array of vectors - corresponding to the output of a convolution layer - and retaining a form of spatial relationship between each vector and a region of the input image (see Fig. 3). Since this spatial relationship is at the heart of this contribution, we purposefully exclude self-explaining image classification models that uses full images as prototypes [5, 6, 8] from our study. Indeed, such architectures use a single latent vector representation, resulting in a trivial relationship between the full image and its latent representation. In this contribution, we rather focus on models where each vector in the latent space corresponds to *a part* of an image. During training, such models extract a set of reference vectors and their visual counterparts from the training set, called *part prototypes* (for simplicity, we simply use the term *prototypes*). Prototypes are either discriminative of a particular class [12, 16, 46] (given by the label of the corresponding training image) or shared among multiple classes [30, 34, 35]. During inference, the similarity between a given prototype and the test image is computed using the L2-distance (or cosine distance [46]) between their respective latent representations, under the assumption that proximity between vectors in the latent space should entail simi-

larity in the visual space. In the case of ProtoPNet [12], ProtoPShare [35], ProtoPool [34], Deformable ProtoPNet [16] and TesNet [46], all similarity scores are then processed through a fully connected layer to produce the final classification. In the case of ProtoTree [30], these similarity scores are used to compute a path across a soft decision tree where each leaf corresponds to one particular class.

It is important to indicate that although our study focuses on ProtoPNet and ProtoTree, all the aforementioned methods (ProtoPShare, ProtoPool, Deformable ProtoPNet, TesNet) share a common code base inherited from ProtoPNet and therefore are theoretically susceptible to the issue raised in this contribution.

Saliency methods Such methods are among the most commonly used building blocks for generating explanations and aim at identifying the most important (salient) features (pixels) of an image w.r.t. the output of a given neuron. Although they are usually used to explain the decision taken by the model (e.g. neurons from the last layer of a classifier), they can also be used to visualise the most salient part of the image w.r.t. any intermediate result of a DCNN. Gradient-based approaches [39] compute the partial derivative $\frac{\delta S}{\delta x_i}$ of the target neuron output S w.r.t. to each input pixel x_i , with improvements such as Integrated Gradients [43] or SmoothGrads [41] aiming at generating more stable explanations. In particular, Smoothgrads "adds noise to remove noise" by averaging gradients over noisy copies $x + n_i$ ($n_i \sim \mathcal{N}(0, \sigma^2)$) of x . Note that in order to take into account the sign and strength of the input [37], it is also possible to perform an element-wise product between the gradient and the input ($\frac{\delta S}{\delta x_i} \odot x$). Layer-wise Relevance Propagation [10] (LRP) uses a set of rules in order to back-propagate the relevance of each individual neuron from the target back to the input, following a conservation property. This method has several variants depending on the applied rules [27]. In the particular context of prototype-based models, [17] proposes a variant of LRP called Prototype Relevance Propagation (PRP), implementing a dedicated rule to propagate relevance across the layer in charge of computing similarity scores, followed by LRP_{COMP} [23]. Interestingly, [4] and [38] have shown that, for DCNNs based on ReLU activation (as it is often the case for the networks used for feature extraction in self-explaining models), Integrated Gradients, LRP with ϵ -rule and Deep-LIFT [37] are all equivalent to $\frac{\delta S}{\delta x_i} \odot X$. For these reasons, in this work, we choose to compare the original part visualisation generated by ProtoPNet and ProtoTree to visualisations generated using Smoothgrads \odot input and PRP. Note that we also exclude Guided-Backpropagation [42] and its application to GRADCAM [36] due to the results of [2] which raise some concerns regarding the relevance of the explanations based on these methods.

We emphasize the fact that a saliency map *is not an explanation by itself*, but is rather used to build an explanation. In the context of self-explaining models based on prototypes, the saliency maps of both the prototype and the test image are cropped in order to retain only the most salient pixels in the corresponding images. The explanation finally consists in the parallel display of the selected regions in the two images.

Evaluation metrics Explanations based on part visualisation implicitly rely on the underlying method used for generating saliency maps. Multiple metrics have been proposed in order to evaluate the quality of explanations against a set of desired properties [29]. In order to answer our first research question (do ProtoPNet and ProtoTree generate faithful image patches corresponding to the latent representations used in their decision-making process?), we focus on the property known as *faithfulness* [3] (a.k.a. *fidelity* [44] or *correctness* [29]) which - in our case - quantifies the adequacy between a saliency method and the model behaviour. In this particular context, faithfulness can primarily be evaluated by model parameter randomisation or deletion/insertion methods. Parameter randomisation [2] analyses the effect of perturbations to the model parameters on the saliency map: if a method produces the same saliency maps regardless of the changes in the model parameters, then it probably does not correctly reflect the model behaviour. As described above, in this regard the results obtained by [2] guide us in our choice of saliency methods. Deletion/insertion metrics [3, 31, 44] monitor the evolution of the target neuron's output when the most/least important pixels of the input image are removed incrementally [31, 44] or individually [3]. Such metrics aim at checking the ability of a saliency method to correctly sort the importance of pixels w.r.t. to a particular model output. In particular, "removing" the most salient pixels identified by a saliency method faithful to the model should result in a strong variation of the target neuron's output, while a less faithful method will highlight pixels whose absence after deletion has actually no influence on the neuron's output. In this work, we choose to implement an incremental deletion metric, limiting ourselves to 2% of the original image in order to avoid evaluating our saliency methods using out-of-distribution inputs [18].

Finally, [28] evaluates the *relevance* of prototype-based explanations by applying controlled perturbations (changes in colour hue, shape, texture, saturation, contrast) on images and monitoring the evolution of the similarity score for each prototype. Note that since these perturbations are applied on the *entire* image, this method is incidentally independent of the spatial location of the target image patch. In this work, we wish to check two properties of ProtoPNet and ProtoTree that are related to our second research question:

after training, does each prototype correspond to a part of the object? during inference, do the patches from the test image that are compared to the prototypes also correspond to parts of the object? Therefore, we propose to focus on the intersection between the pixels highlighted by the saliency methods and the ground-truth segmentation of the object (when available). We do not use the popular Intersection-over-Union (IoU) metric because our goal is not to assert whether each prototype covers the entire object, but rather measure the percentage of the image patch intersecting the object segmentation.

3. Theoretical background

In this work, we consider a classification problem with a training set $X_{train} \subset \mathcal{X} \times \mathcal{Y}$. Let f be a fully convolutional neural network (fCNN) processing images in \mathcal{X} and producing a D -dimensional latent representation of size (H, W) . We denote $\mathcal{L}_f = \mathbb{R}^{H \times W \times D}$ the latent space associated with f . For $x \in \mathcal{X}$, we denote $f^{(h,w)}(x) \in \mathbb{R}^D$ the D -dimensional vector corresponding to the h -th row and w -th column of $f(x)$.

3.1. ProtoPNet and ProtoTree

Finding similarities in the latent space For $x \in \mathcal{X}$, ProtoPNet and ProtoTree compute their decision (classification) based on similarities between the latent representation $f(x)$ and a set of D -dimensional vectors (r_1, \dots, r_p) that are learned during training and act as reference points in the latent space. More precisely, the similarity between r_i and a particular vector $f^{(h,w)}(x)$ is defined as $s_i^{(h,w)}(x) = \log((\|f^{(h,w)}(x) - r_i\|_2^2 + 1) / (\|f^{(h,w)}(x) - r_i\|_2^2 + \epsilon))$ (ProtoPNet) or $s_i^{(h,w)}(x) = e^{-\|f^{(h,w)}(x) - r_i\|_2^2}$ (ProtoTree), where $\|\cdot\|_2$ denotes the L2 distance. For each reference point r_i , $s_i(x) \in \mathbb{R}^{H \times W}$ is called the *similarity map* between x and r_i . The model decision process d is a function of the aggregation of high similarity scores $s(x) = (\max(s_1(x)), \dots, \max(s_p(x)))$: weighted sum for ProtoPNet, soft decision tree for ProtoTree. During training, the parameters (convolutional weights) of the feature extractor f , of the decision function¹ d , and the reference points r_i are jointly learned in order to minimize the cross-entropy loss between the prediction $d \circ s(x) \in \mathcal{Y}$ and y , $\forall (x, y) \in X_{train}$.

Prototype projection After training, the reference points r_i are "pushed" toward latent representations of parts of training images, in a process called *prototype projection*. More formally, a prototype P_i is a tuple $P_i = (p_i, h_i, w_i, r_i)$, computed using a projection dataset

$X_{proj} \subset \mathcal{X}$ (usually consisting of images from the training set X_{train}), where

$$\begin{cases} p_i, h_i, w_i = \arg \min_{x \in X_{proj}, h, w} \|f^{(h,w)}(x) - r_i\|_2^2 \\ r_i \leftarrow f^{(h_i, w_i)}(p_i) \end{cases} \quad (1)$$

Thus, prototype projection moves each reference point - corresponding to an abstract vector in the latent space - to a nearby point which is, by construction, obtained from an image p_i . Although this operation may reduce the accuracy of the system (changing the value of the reference points r_i may increase the cross-entropy loss), it allows the system to anchor the reference points as latent representations of real images from the projection set.

From similarity map to part visualisation Given an image x and a prototype P_i , ProtoPNet generates a visualisation of the most similar image patch in x by upsampling the similarity map $s_i(x) \in \mathbb{R}^{H \times W}$ to the size of x using cubic interpolation, then cropping the resulting saliency map to the 95th percentile, while ProtoTree retains only the location of highest similarity in $s_i(x)$ before upsampling (setting all other locations to 0), as shown in Fig. 2. Note that the same method is also used to visualise the prototype itself by setting $x = p_i$ and that prototype visualisation can be performed once and for all after projection. However, both approaches do not factor in the size of the receptive field² of each neuron in the last layers of a deep CNN. Indeed, taking into account padding and pooling layers in the architecture of DCNNs, the value $s_i^{(h,w)}(x)$ may actually depend on the entire image x [7] rather than a localized region.

This looks like that As illustrated in Fig. 3, during inference, for an image $x \in \mathcal{X}$, both ProtoPNet and ProtoTree base their decision on the comparison between the latent representation $f(x)$ and the representatives of the prototypes. More precisely, for each prototype P_i , the model finds the vector in $f(x)$ closest to r_i , corresponding to the highest score of the similarity map $s_i(x)$. If this score is above a given threshold, then it generates and shows side by side patches of images extracted from the prototype image p_i and the test image x .

3.2. Generating part visualisation with PRP and Smoothgrads

Similar to ProtoTree, for an image $x \in \mathcal{X}$ and a prototype P_i , we first find the coordinates and value of the highest

¹The details of the decision function d are not relevant to our work, which focuses on the method used by ProtoPNet and ProtoTree to build a saliency map out of the similarity map $s_i(x)$.

²This information is actually computed in the code of ProtoPNet, but never put to use.

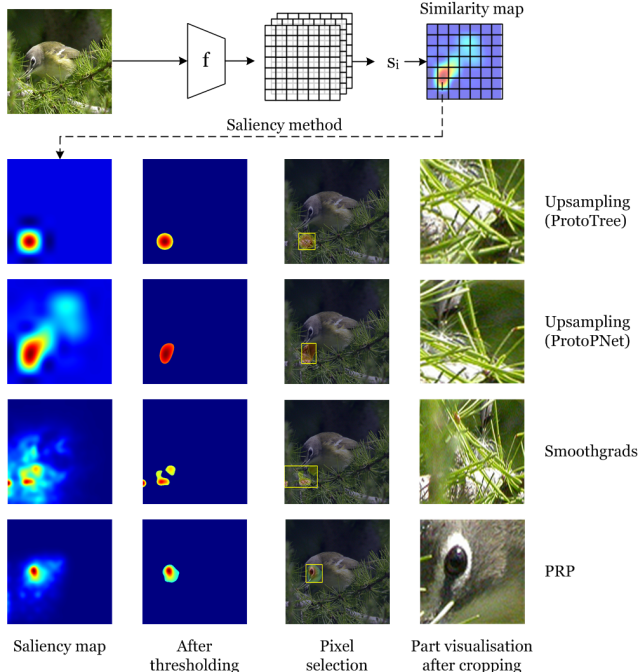


Figure 2. Generating part visualisation from the similarity map. First, a saliency method generates a saliency map from the similarity map. Then, thresholding is applied to retain only the most salient pixels. Finally the original image is cropped to produce a part visualisation.

similarity score

$$\begin{cases} h_m, w_m = \arg \max_{h,w} s_i^{(h,w)}(x) \\ s_m = \max_{h,w} s_i^{(h,w)}(x) = s_i^{(h_m, w_m)}(x) \end{cases} \quad (2)$$

Where ProtoPNet and ProtoTree directly upsample the similarity map to the size of the original image x , in this work we generate saliency maps identifying the most important pixels w.r.t. the highest similarity score s_m by applying Smoothgrads or PRP on the output of the neuron $s_i^{(h_m, w_m)}$. Importantly, we perform a back-propagation of the similarity score *through* the feature extractor f in order to take into account its architecture and parameters. Then, we obtain a part visualisation by using the saliency map to retain only the 2% of most important pixels from x and cropping the image accordingly, as shown in Fig. 2. Again, the same method is applied in order to extract a part visualisation for each prototype.

3.3. Measuring faithfulness

As illustrated in Fig. 4, and similar to [31], we measure the faithfulness of a saliency method to the model behaviour by analysing the drop in similarity when incrementally “deleting” pixels of highest relevance. In practice,

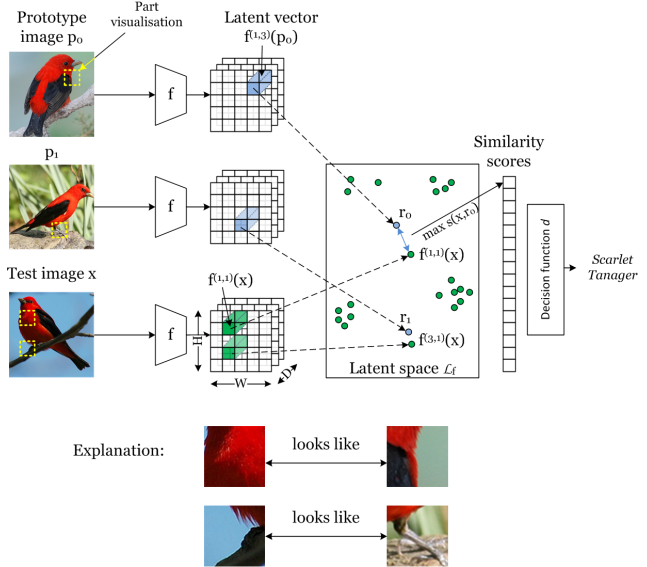


Figure 3. ProtoPNet/ProtoTree inference overview. After training and projection, a prototype P_i is composed of an image from the training set (p_i) and a latent vector (e.g. $r_0 = f^{(1,3)}(p_0)$). During inference, similarity scores are computed based on the closest distance to each prototype latent representation and used in the decision function d to produce the prediction. Then, part visualisation is applied to associate each latent vector to a patch inside of its original image. Finally, an explanation is generated by visualising side by side all relevant prototypes and their most similar image patches in x . Best viewed in colour.

for an image $x \in \mathcal{X}$, a prototype P_i and a given saliency method (upsampling, Smoothgrads, PRP), we first generate the saliency map from the similarity map $s_i(x)$. Then, for a target deletion area a , we measure the drop in similarity as follows:

- We generate an image \tilde{x} obtained after masking out the $a\%$ most salient pixels from x (see Fig. 5).
- We compute the value $\tilde{s}_m = s^{(h_m, w_m)}(\tilde{x})$ corresponding to the similarity score of \tilde{x} at the original location of highest similarity.
- We compute the value $\tau(a) = \tilde{s}_m / s_m$ measuring the relative drop in similarity between x and \tilde{x} .

$\tau(a) \approx 1$ indicates that deleting these pixels has no impact on the similarity score, thus that the saliency method is not faithful to the model. On the contrary, $\tau(a) \approx 0$ indicates that deleting these pixels has a high impact on the similarity score, thus that the saliency method correctly identifies relevant pixels w.r.t. to the similarity score and is faithful to the model. Finally, we compare the faithfulness of all saliency methods under test by computing the Area under the Deletion Curve (AUDC) for values of a up to 2% of the

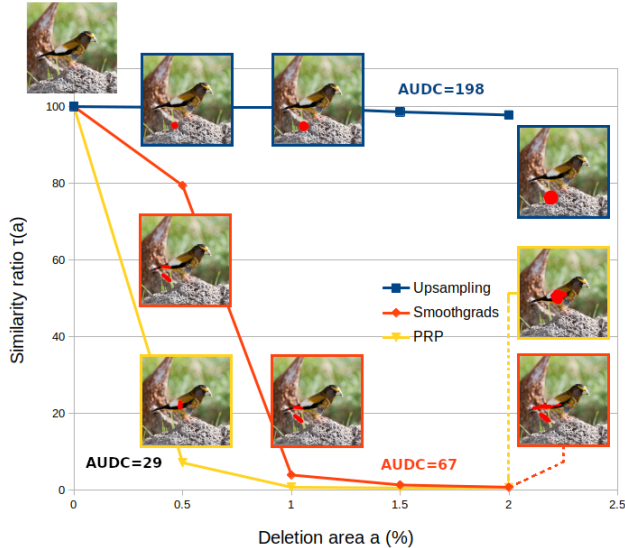


Figure 4. Evolution of the similarity ratio when incrementally removing the most important pixels according to the ProtoPNet/ProtoTree method (upsampling, in blue), Smoothgrads (red), and PRP (yellow). In this example, removing pixels according to upsampling has little to no effect on the similarity score, suggesting that the explanation is incorrect. On the contrary, when removing only 1% of the image according to Smoothgrads, the similarity score drops to 2% of its original value, suggesting that the explanation focuses on actual regions of interest for the model. The same result is achieved when removing only 0.5% of the image according to PRP, indicating an even more precise explanation. Moreover, reaching a similarity ratio lower than 10% indicates that the explanation method has successfully identified the most relevant pixels of the image patch and gives an indication on the effective size of the image patch. Best viewed in colour.

total image area (as indicated in Sec. 2, we restrict ourselves to deleting a small portion of the original image in order to avoid unexpected behaviours from the DCNNs [18]). Interestingly, identifying the minimum deletion area leading to a significant drop in similarity also provides an indication regarding the size of the **effective receptive field** of the model w.r.t. to a given neuron. Indeed, as shown in Fig. 4, reaching a similarity ratio $\tau(a) < 0.1$ indicates that the deleted pixels amount to 90% of the similarity score, thus that the area of the effective receptive field is close to a .

3.4. Measuring relevance

In order to verify that the model is producing decisions based on parts of the object rather than elements of the background, we measure the relevance of both prototypes and their most similar patches in test images. As mentioned in Sec. 2, we measure the relevance of an image patch as its intersection with the object segmentation, assuming that such information is available in the training data. In practice, for a given saliency method, an input image x and a prototype

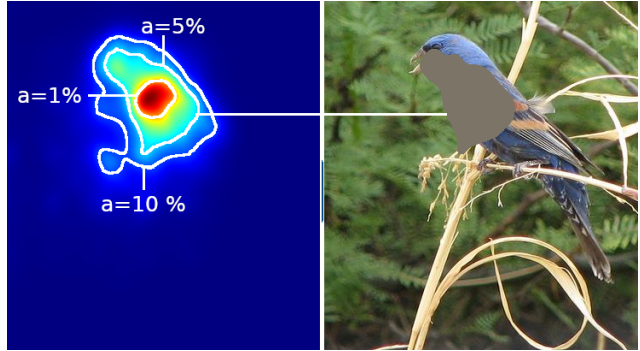


Figure 5. From a saliency map generated by a saliency method (left), we progressively mask out a growing area of image x , selecting the most salient pixels first, and generate perturbed images \tilde{x} (right). By studying the impact of perturbed images on the similarity score, the deletion metric evaluates the faithfulness of the saliency method.

P_i , we first identifies the 2% most salient pixels of x w.r.t. the maximum similarity score. If few or none of these pixels (less than 5% in practice) intersect the object segmentation, then the image patch is considered irrelevant, as it mostly focuses on the background. As shown in Fig. 1, different methods lead to different saliency maps that produce image patches with different relevance. The effective relevance of prototypes and patches from the test image is therefore dependent on the faithfulness of the chosen saliency method.

4. Experiments

Setup We perform our experiments on two popular fine-grained datasets: the CUB-200-2011 [47] dataset (CUB) contains 11,788 images belonging to 200 bird species - split into 5,994 training images and 5,794 test images - and provides the object bounding box coordinates and segmentation mask for each image; the Stanford Cars [49] dataset (CARS) contains 16,185 images belonging to 196 car models, evenly split into training and test images. Additionally, for ProtoPNet, we use the cropped images of the CUB dataset - which we denote CUB-c - during training and inference. As summarized in Table 1, for both models, we primarily use a Resnet50 [19] feature extractor (*backbone*), pretrained on the iNaturalist [21] dataset (CUB) or the ImageNet [14] dataset (CARS), with images of size 224×224 . In order to compare results across different feature extractors, we also train a ProtoPNet on CUB using a VGG19 [40] network (pretrained on ImageNet).

For saliency methods, we use the code of PRP [17] kindly provided by the authors. For Smoothgrads [41], we use 10 noisy samples per image and dynamically set the value of σ using a noise ratio of 0.2 (*i.e.* $\sigma = 0.2 \times \Delta x$). For both methods, we post-process the saliency map by successively averaging the gradients at each pixel location across

Model	Backbone	Dataset	Accuracy
ProtoPNet	VGG19	CUB-c	75.1%
	ResNet50	CUB-c	72.5%
		CARS	71.4%
ProtoTree	ResNet50	CUB	83.1%
		CARS	83.2%

Table 1. Accuracy of the self-explaining models used in this work. CUB-c denotes the cropped CUB-200-2011 dataset.

Model	Backbone	Dataset	Up.	Method PRP	$S \odot I$
Prototypes					
ProtoPNet	VGG19	CUB-c	82/148	77/139	74/136
	ResNet50	CUB-c	78/142	62/121	74/132
		CARS	91/176	61/136	68/141
ProtoTree	ResNet50	CUB	196/190	153/130	181/164
		CARS	192/181	150/142	175/168

Table 2. Average Area Under the Deletion Curve (AUDC \downarrow) of prototypes (left value) and test patches (right value) generated by ProtoPNet and Prototree when using upsampling, Smoothgrads ($S \odot I$) and PRP. For readability, all values are rescaled by a factor 10,000, with AUDC=200 corresponding to a similarity ratio τ maintained at 100% for all deletion areas up to 2% (see Fig. 4).

the RGB channels, taking the absolute value (putting equal emphasis on strongly positive and negative gradients), and applying a 5×5 Gaussian filter in order to avoid isolated gradients due to max-pooling layers inside of f .

Faithfulness of patch visualisation We evaluate the faithfulness of the saliency methods under test by measuring the AUDC (as described in Sec. 3.3) when visualising prototypes after training, but also on patches of images from the test set during inference: for ProtoTree, we apply the saliency method only when a positive comparison is established between a patch of the test image and a prototype (right branch of each decision node); for ProtoPNet, we apply the saliency method on the 10 patches of the test image that are most similar to any prototype of the inferred class. The AUDC score is measured by using deletion areas between 0% and 2%, with an increment value of 0.1%. As shown in Table 2, in all cases, the upsampling method used in ProtoPNet and ProtoTree leads to a higher AUDC score than Smoothgrads or PRP. This confirms the imprecision pointed out in [17] and extends the issue to ProtoTree. Moreover, except in the case of ProtoPNet on CUB-c using a VGG19 backbone, PRP seems to provide a more faithful saliency maps than Smoothgrads (lower AUDC), at the cost of a higher computing time due to its rules for relevance propagation (on a Nvidia Quadro T2000 Mobile, a single PRP relevance propagation is approximately 1.4 times slower than Smoothgrads with 10 noisy samples). However, we note that the AUDC scores between upsampling, Smoothgrads and PRP are fairly similar on the CUB-

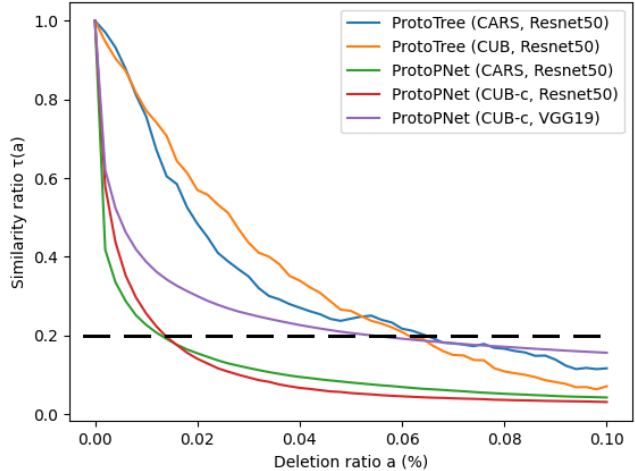


Figure 6. Average similarity ratio v. deletion area when using PRP for visualising prototypes. The average similarity ratio drops "faster" for ProtoPNet prototypes than for ProtoTree, suggesting a greater effective receptive field for ProtoTree prototypes.

c dataset, which is probably due to the image cropping that increases the relative area of the bird inside of the image and decreases the probability to miss the important pixels, even with a random guess.

Interestingly, as shown in Fig. 6, when extending the deletion area to 10% of the image, we note that on average, the drop in similarity ratio ($\tau(a) < 0.2$) occurs below 2% for ProtoPNet prototypes when using ResNet50, and around 7% for ProtoTree prototypes or ProtoPNet using VGG19. Since the similarity ratio eventually reaches values below 0.2, this result does not question the faithfulness of PRP explanations. As mentioned in Sec. 3.3, this effect rather suggests that the size of the effective receptive field of ProtoTree prototypes (or ProtoPNet using VGG19) is larger in general than ProtoPNet with ResNet50, even on the same dataset (CARS). For ProtoPNet with VGG19, this suggests a sensitivity of the model to the underlying backbone, an hypothesis that is reinforced by the next experiment. For ProtoTree, this may be due to the fact that the decision tree shares prototypes among all classes and therefore does not focus on very small details, contrary to ProtoPNet. This effect is also present when using Smoothgrads but seems uncorrelated to the depth of the prototype inside the decision tree (see Supplementary material). Moreover, this clarifies the discrepancy in AUDC scores between ProtoPNet and ProtoTree visualisations. In particular, using the same feature extractor (Resnet50) on the same dataset (CARS), ProtoPNet visualisation with PRP reaches significantly lower AUDC scores (61 for prototypes, 136 for test patches) than ProtoTree (150 for prototypes, 142 for test patches).

This first experiment confirms that the upsampling method implemented in ProtoPNet and ProtoTree produces less

Model	Backbone	Up.	Method	
			PRP	$S \odot I$
ProtoPNet	VGG19	15.4% / 23.3%	10.% / 16.6%	11.35% / 19.9%
	ResNet50	2.0% / 8.8%	1.3% / 8.0%	1.0% / 6.1%
ProtoTree	ResNet50	35.4% / 51.9%	0.5% / 0.5%	8.7% / 14.5%

Table 3. Percentage of prototypes (left value) and test patches (right value) with less than 5% of intersection with the object (CUB dataset).

faithful image patches than PRP of Smoothgrads. Therefore, the explanations provided by default in these models do not necessarily reflect the actual behaviour of the model.

Relevance of patch visualisation In this experiment, we use the segmentation masks provided by the CUB-200-2011 dataset (such information is not provided with the Stanford Cars dataset). As indicated in Sec. 3.4, we measure the percentage of saliency masks not intersecting the object (intersection of less than 5%), for both prototypes and test image patches. As shown in Table 3 and illustrated in Fig. 7, the imprecision of the upsampling visualisation method used in ProtoPNet and ProtoTree leads to a false sense of model bias. In particular, when using the upsampling method, more than a third (35.4%) of ProtoTree prototypes and half (51.9%) of the test image patches seem to be focusing on elements of the background rather than the bird. However, when using a more faithful saliency method such as PRP (or even Smoothgrads), we notice that only 0.5% of the prototypes (*i.e.* a single prototype) and 0.5% of test image patches are actually irrelevant. Note that this gap between results is again more limited for ProtoPNet where images from the CUB dataset are cropped (in order to achieve a better accuracy) and where the upsampling method is less likely to "miss" the object entirely. Finally, when applying the PRP method on ProtoPNet, we also notice that the percentage of biased prototypes and test image patches is significantly more important when using VGG19 as a backbone, compared to using Resnet50, which raises the question of the sensitivity of prototype-based architectures to the underlying backbone architecture or to the pre-trained features.

In this second experiment, we have shown that the apparent bias of ProtoPNet and ProtoTree suggested by the use of upsampling is largely mitigated when using PRP and Smoothgrads. Far from contradicting the results of [17], this mainly confirms that the use of an unfaithful saliency method for generating image patches provide unreliable information regarding the actual decision-making process of the model and can be detrimental to the trust in self-explaining models in general.

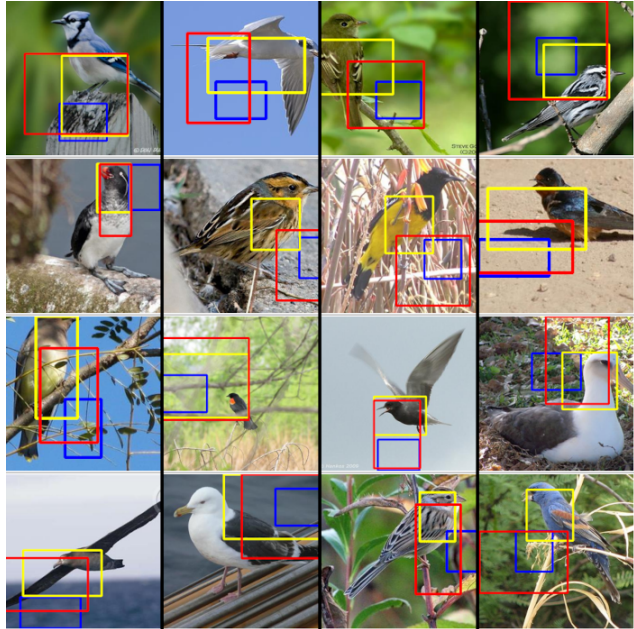


Figure 7. Visualisation of prototypes from a ProtoTree trained on CUB-200-2011 using upsampling with cubic interpolation (blue), Smoothgrads (red) or PRP (yellow). In these examples, the upsampling strategy misses the object, given a false sense of bias in the model. Best viewed in colour.

5. Discussion and future work

Although case-based reasoning architectures for image classification constitute a first stepping stone towards more interpretable computer vision models, such architectures still suffer from several shortcomings that may hinder their widespread usage, especially in critical applications. In this work, we have shown that even though such models might produce a correct decision for the right reasons (*this indeed looks like that*), they may yet fail to properly explain this decision by incorrectly identifying appropriate parts of the images (prototypes and test image patches). In particular, more faithful saliency methods such as PRP and Smoothgrads can help uncover biases [17] or - in our experiments - disprove *apparent* biases of the model. As indicated in Sec. 2, this issue is likely not restricted to ProtoPNet or ProtoTree, since ProtoPShare, ProtoPool, Deformable ProtoPNet and TesNet share a common code base inherited from ProtoPNet that includes the upsampling method for patch visualisation.

Additionally, in order to achieve performance (*e.g.* accuracy) on par with non-interpretable methods, case-based reasoning architectures rely on complex feature extractors such as DCNNs, under the assumption that proximity in the latent space entails similarity in the visual space. However, as shown by [20], such assumption may not always hold. Moreover, according to our own experiments, the relevance

of prototypes and test image patches may vary depending on the choice of backbone architecture (*e.g.* Resnet50 or VGG19), or at least depending on the nature of the pre-trained features used at the beginning of the training process. Finally, proving that the model is indeed focusing on the object to produce its decision does not imply that such decision is based on understandable rather than abstract information. In this regard, although the work of [28] can help shed some light on the visual cues (colour, shape, hue, etc) used to determine similarity, some similarities between image patches remain unintelligible. This suggests that case-based reasoning architectures using DCNNs for feature extraction are not actually explainable-*by-design*, in the sense that a decision-making process based on distance in the latent space is not sufficient to guarantee interpretability. As a consequence, we also argue that raw performance (*e.g.* classification accuracy) should not be used to compare such architectures, as it may drive the research community towards models focusing on more abstract features.

In a future work, we first wish to extend our study to other case-based reasoning models. In particular, ProtoPool [34] uses a focal similarity to ensure the locality of prototypes and to reduce the probability of the model learning elements of the background. Using the metric described in Sec. 3.4, this gain in prototype relevance could be quantified. Finally, as stated above, there exists a dire need for metrics for evaluating the understandability of explanations in a systematic manner and, in the case of prototype-based architectures, to properly quantify visual similarity.

References

- [1] Amina Adadi and Mohammed Berrada. Peeking inside the black-box: A survey on explainable artificial intelligence (xai). *IEEE Access*, 6:52138–52160, 2018. 1
- [2] Julius Adebayo, Justin Gilmer, Michael Muelly, Ian Goodfellow, Moritz Hardt, and Been Kim. Sanity Checks for Saliency Maps. In *Advances in Neural Information Processing Systems* 32, page 11, 2018. 3
- [3] David Alvarez-Melis and T. Jaakkola. Towards robust interpretability with self-explaining neural networks. *ArXiv*, abs/1806.07538, 2018. 1, 3
- [4] Marco Ancona, Enea Ceolini, A. Cengiz Öztireli, and Markus H. Gross. A unified view of gradient-based attribution methods for deep neural networks. *ArXiv*, abs/1711.06104, 2017. 3
- [5] Plamen P. Angelov and Eduardo A. Soares. Towards explainable deep neural networks (xdnn). *Neural networks : the official journal of the International Neural Network Society*, 130:185–194, 2019. 2
- [6] Plamen P. Angelov and Eduardo A. Soares. Towards deep machine reasoning: a prototype-based deep neural network with decision tree inference. *2020 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, pages 2092–2099, 2020. 2
- [7] André Araujo, Wade Norris, and Jack Sim. Computing receptive fields of convolutional neural networks. *Distill*, 2019. <https://distill.pub/2019/computing-receptive-fields>. 4
- [8] Sercan Ö. Arik and Tomas Pfister. Protoattend: Attention-based prototypical learning. *J. Mach. Learn. Res.*, 21:210:1–210:35, 2019. 2
- [9] Alejandro Barredo Arrieta, Natalia Díaz Rodríguez, Javier Del Ser, Adrien Bénézet, Siham Tabik, A. Barbado, Salvador García, Sergio Gil-Lopez, Daniel Molina, Richard Benjamins, Raja Chatila, and Francisco Herrera. Explainable artificial intelligence (xai): Concepts, taxonomies, opportunities and challenges toward responsible ai. *Inf. Fusion*, 58:82–115, 2019. 1
- [10] Alexander Binder, Sebastian Bach, Gregoire Montavon, Klaus-Robert Müller, and Wojciech Samek. Layer-wise relevance propagation for deep neural network architectures. In Kuinam J. Kim and Nikolai Joukov, editors, *Information Science and Applications (ICISA) 2016*, pages 913–922, Singapore, 2016. Springer Singapore. 1, 3
- [11] Francesco Bodria, Fosca Giannotti, Riccardo Guidotti, Francesca Naretto, Dino Pedreschi, and Salvatore Rinzivillo. Benchmarking and survey of explanation methods for black box models. *ArXiv*, abs/2102.13076, 2021. 1
- [12] Chaofan Chen, Oscar Li, Chaofan Tao, Alina Jade Barnett, Jonathan Su, and Cynthia Rudin. This looks like That: Deep learning for interpretable image recognition. *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, page 8930–8941, 2019. 1, 2, 3
- [13] Zhi Chen, Yijie Bei, and Cynthia Rudin. Concept whitening for interpretable image recognition. *Nature Machine Intelligence*, 2(12):772–782, Dec. 2020. 1
- [14] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image

- database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 6
- [15] Amit Dhurandhar, Tejaswini Pedapati, Avinash Balakrishnan, Pin-Yu Chen, Karthikeyan Shanmugam, and Ruchi Puri. Model agnostic contrastive explanations for structured data. *ArXiv*, abs/1906.00117, 2019. 1
- [16] Jonathan Donnelly, Alina Jade Barnett, and Chaofan Chen. Deformable protopnet: An interpretable image classifier using deformable prototypes. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10255–10265, 2021. 1, 2, 3
- [17] Srishti Gautam, Marina M.-C. Höhne, Stine Hansen, Robert Jensen, and Michael Kampffmeyer. This looks more like that: Enhancing self-explaining models by prototypical relevance propagation. *Pattern Recognition*, page 109172, 2022. 2, 3, 6, 7, 8
- [18] Tristan Gomez, Thomas Fr’eur, and Harold Mouchère. Metrics for saliency map evaluation of deep learning explanation methods. In *International Conferences on Pattern Recognition and Artificial Intelligence*, 2022. 3, 6
- [19] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016. 2, 6
- [20] Adrian Hoffmann, Claudio Fanconi, Rahul Rade, and Jonas Kohler. This looks like that... does it? shortcomings of latent space prototype interpretability in deep networks. 2021. 8
- [21] Grant Van Horn, Oisin Mac Aodha, Yang Song, Yin Cui, Chen Sun, Alexander Shepard, Hartwig Adam, Pietro Perona, and Serge J. Belongie. The inaturalist species classification and detection dataset. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8769–8778, 2017. 6
- [22] Been Kim, Martin Wattenberg, Justin Gilmer, Carrie J. Cai, James Wexler, Fernanda B. Viégas, and Rory Sayres. Interpretability beyond feature attribution: Quantitative testing with concept activation vectors (tcav). In *International Conference on Machine Learning*, 2017. 1
- [23] Maximilian Kohlbrenner, Alexander Bauer, Shinichi Nakajima, Alexander Binder, Wojciech Samek, and Sebastian Lapuschkin. Towards best practice in explaining neural network decisions with lrp. *2020 International Joint Conference on Neural Networks (IJCNN)*, pages 1–7, 2019. 3
- [24] Thibault Laugel, Marie-Jeanne Lesot, Christophe Marsala, X. Renard, and Marcin Detryniecki. Comparison-based inverse classification for interpretability in machine learning. In *International Conference on Information Processing and Management of Uncertainty*, 2018. 1
- [25] Arnaud Van Looveren and Janis Klaise. Interpretable counterfactual explanations guided by prototypes. *ArXiv*, abs/1907.02584, 2019. 1
- [26] Scott M. Lundberg and Su-In Lee. A unified approach to interpreting model predictions. In *NIPS*, 2017. 1
- [27] Grégoire Montavon, Alexander Binder, Sebastian Lapuschkin, Wojciech Samek, and Klaus-Robert Müller. Layer-wise relevance propagation: An overview. In *Explainable AI*, 2019. 3
- [28] Meike Nauta, Annemarie Jutte, Jesper C. Provoost, and Christin Seifert. This looks like that, because ... explaining prototypes for interpretable image recognition. In *PKDD/ECML Workshops*, 2020. 3, 9
- [29] Meike Nauta, Jan Trienes, Shreyasi Pathak, Elisa Nguyen, Michelle Peters, Yasmin Schmitt, Jörg Schlötterer, Maurice van Keulen, and Christin Seifert. From anecdotal evidence to quantitative evaluation methods: A systematic review on evaluating explainable AI. *CoRR*, abs/2201.08164, 2022. 1, 3
- [30] Meike Nauta, Ron van Bree, and Christin Seifert. Neural prototype trees for interpretable fine-grained image recognition. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14928–14938, 2021. 1, 2, 3
- [31] Vitali Petsiuk, Abir Das, and Kate Saenko. Rise: Randomized input sampling for explanation of black-box models. In *BMVC*, 2018. 3, 5
- [32] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. ”why should i trust you?”: Explaining the predictions of any classifier. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2016. 1
- [33] Cynthia Rudin. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1:206–215, 2019. 1
- [34] Dawid Rymarczyk, Lukasz Struski, Michal G’orszczak, Koryna Lewandowska, Jacek Tabor, and Bartosz Zieliński. Interpretable image classification with differentiable prototypes assignment. In *European Conference on Computer Vision*, 2021. 1, 2, 3, 9
- [35] Dawid Rymarczyk, Lukasz Struski, Jacek Tabor, and Bartosz Zieliński. Protoshare: Prototypical parts sharing for similarity discovery in interpretable image classification. *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, 2021. 1, 2, 3
- [36] Ramprasaath R. Selvaraju, Abhishek Das, Ramakrishna Vedantam, Michael Cogswell, Devi Parikh, and Dhruv Batra. Grad-cam: Why did you say that? *ArXiv*, abs/1611.07450, 2016. 3
- [37] Avanti Shrikumar, Peyton Greenside, and Anshul Kundaje. Learning important features through propagating activation differences. In *International Conference on Machine Learning*, 2017. 3
- [38] Avanti Shrikumar, Peyton Greenside, Anna Shcherbina, and Anshul Kundaje. Not just a black box: Learning important features through propagating activation differences. *ArXiv*, abs/1605.01713, 2016. 1, 3
- [39] Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Deep inside convolutional networks: Visualising image classification models and saliency maps. *CoRR*, abs/1312.6034, 2013. 1, 3
- [40] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *CoRR*, abs/1409.1556, 2015. 2, 6
- [41] Daniel Smilkov, Nikhil Thorat, Been Kim, Fernanda B. Viégas, and Martin Wattenberg. Smoothgrad: removing

- noise by adding noise. *ArXiv*, abs/1706.03825, 2017. [1](#), [2](#), [3](#), [6](#)
- [42] Jost Tobias Springenberg, Alexey Dosovitskiy, Thomas Brox, and Martin A. Riedmiller. Striving for simplicity: The all convolutional net. *CoRR*, abs/1412.6806, 2014. [3](#)
- [43] Mukund Sundararajan, Ankur Taly, and Qiqi Yan. Axiomatic attribution for deep networks. In Doina Precup and Yee Whye Teh, editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 3319–3328. PMLR, 06–11 Aug 2017. [1](#), [3](#)
- [44] Richard J. Tomsett, Daniel Harborne, Supriyo Chakraborty, Prudhvi K. Gurram, and Alun David Preece. Sanity checks for saliency metrics. *ArXiv*, abs/1912.01451, 2019. [3](#)
- [45] Sandra Wachter, Brent Daniel Mittelstadt, and Chris Russell. Counterfactual explanations without opening the black box: Automated decisions and the gdpr. *Cybersecurity*, 2017. [1](#)
- [46] Jiaqi Wang, Huafeng Liu, Xinyue Wang, and Liping Jing. Interpretable image recognition by constructing transparent embedding space. *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 875–884, 2021. [2](#), [3](#)
- [47] P. Welinder, S. Branson, T. Mita, C. Wah, F. Schroff, S. Belongie, and P. Perona. Caltech-UCSD Birds 200. Technical Report CNS-TR-2010-001, 2010. [1](#), [2](#), [6](#)
- [48] Weibin Wu, Yuxin Su, Xixian Chen, Shenglin Zhao, Irwin King, Michael R. Lyu, and Yu-Wing Tai. Towards global explanations of convolutional neural networks with concept attribution. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8649–8658, 2020. [1](#)
- [49] L. Yang, Ping Luo, Chen Change Loy, and Xiaoou Tang. A Large-Scale Car Dataset for Fine-Grained Categorization and Verification. *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3973–3981, 2015. [1](#), [2](#), [6](#)

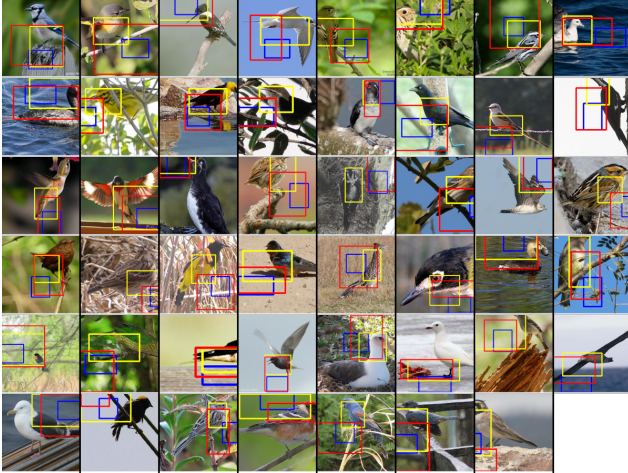


Figure 8. More examples of visualization of prototypes from a ProtoTree trained on CUB-200-2011 using upsampling with cubic interpolation (blue), Smoothgrads (red) or PRP (yellow).

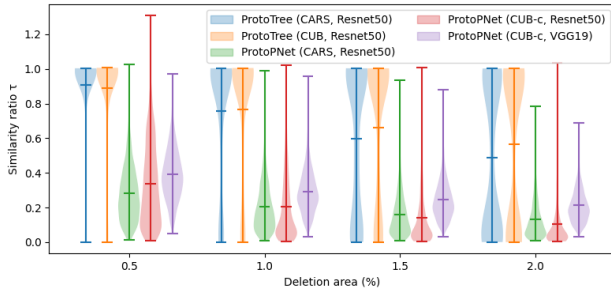


Figure 9. Distributions of similarity ratios v. percentage of deletion area when visualising prototypes using PRP.

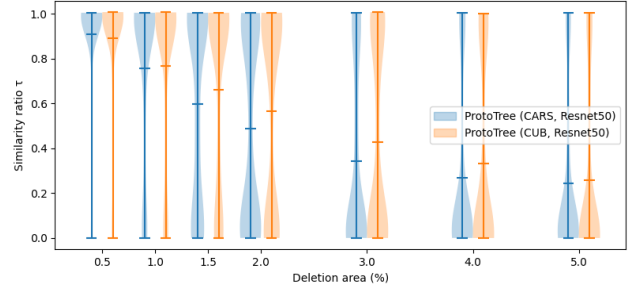
6. Supplementary material

6.1. More prototype visualisation with Smoothgrads and PRP

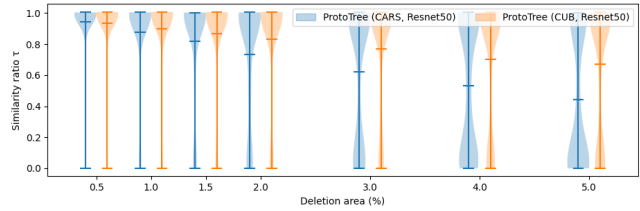
Fig. 8 illustrates how using more faithful visualisation methods, such as PRP or Smoothgrads, rather than upsampling can improve the trust that the user can have in the model. In these examples, the upsampling strategy shows image patches focused on the background and gives a false sense of bias in the model, while PRP and Smoothgrads - which provide more faithful saliency maps - are focusing on elements of the bird.

6.2. Distribution of similarity ratio v. deletion area on ProtoTree visualisation

In addition to the results presented in the paper focusing on the average similarity ratio v. deletion area, in this section we study the *distribution* of similarity ratios for a given deletion area (here 0.5%, 1%, 1.5% and 2%). As shown in Fig. 9, we notice a "sandglass" effect on the distribution of

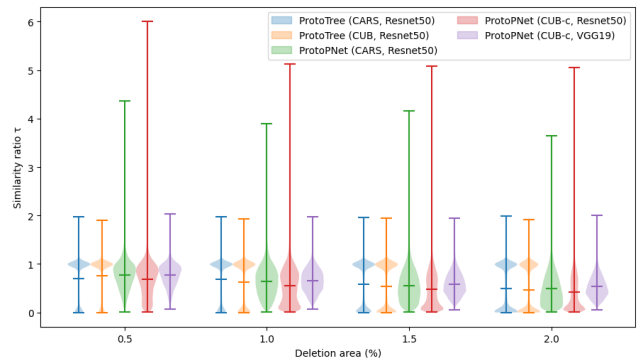


(a) Using PRP.

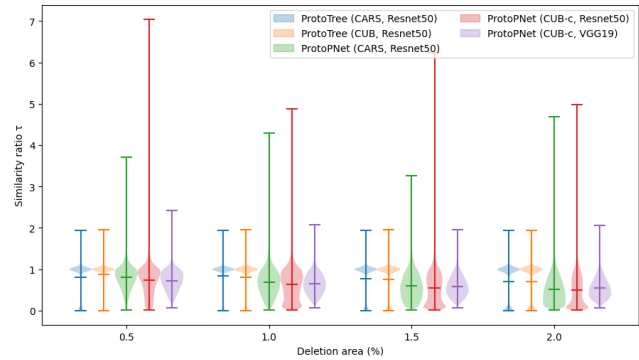


(b) Using Smoothgrads.

Figure 10. Distributions of similarity ratios v. percentage of deletion area when visualising ProtoTree prototypes



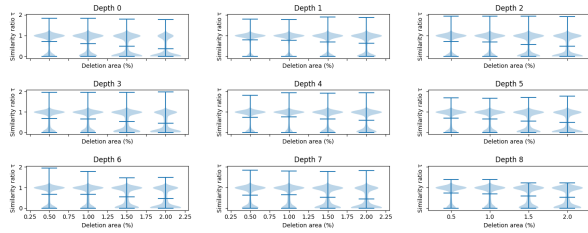
(a) Using PRP



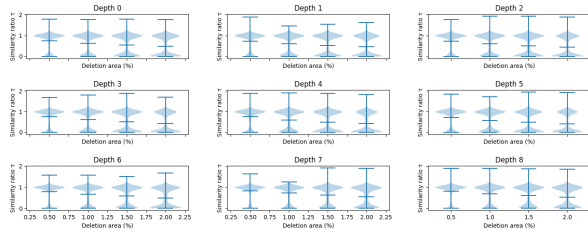
(b) Using Smoothgrads

Figure 11. Distributions of similarity ratios v. percentage of deletion area when visualising test patches during inference.

similarity ratios for ProtoTree prototypes: for low deletion areas ($\leq 1\%$), the similarity ratio for all prototypes is close to 1. Then, from 1.5% up to 4-5% (Fig. 10a), the distribution



(a) On CARS



(b) On CUB

Figure 12. Distributions of similarity ratios v. percentage of deletion area when visualising ProtoTree test patches using PRP during inference. Results are sorted by depth inside of the decision tree.

of similarity ratios slowly shifts towards 0. This suggests that the drop in similarity does occur uniformly for all prototypes, but rather in a "continuous" manner, *i.e.* that ProtoTree prototypes have a wider range of size for their corresponding effective receptive fields than ProtoPNet prototypes. Moreover, as shown in Fig. 10b, the sandglass effect is also present when using Smoothgrads and when visualising image patches during inference (Fig. 11). Moreover, as shown in Fig. 12, this effect is seemingly uncorrelated to the depth of the prototype inside of the decision tree. This suggests that ProtoTree does not necessarily focus on finer - and smaller - details in the last stages of the decision process.