



OPEN

On the interpretability of part-prototype based classifiers: a human centric analysis

Omid Davoudi¹✉, Shayan Mohammadzadehsamakosh² & Majid Komeili¹

Part-prototype networks have recently become methods of interest as an interpretable alternative to many of the current black-box image classifiers. However, the interpretability of these methods from the perspective of human users has not been sufficiently explored. In addition, previous works have had major issues with following proper experiment design and task representation that limit their reliability and validity. In this work, we have devised a framework for evaluating the interpretability of part-prototype-based models from a human perspective that solves these issues. The proposed framework consists of three actionable metrics and experiments. The results of these experiments will reveal important and reliable interpretability related properties of such models. To demonstrate the usefulness of our framework, we performed an extensive set of experiments using Amazon Mechanical Turk. They not only show the capability of our framework in assessing the interpretability of various part-prototype-based models, but they also are, to the best of our knowledge, the most comprehensive work on evaluating such methods in a unified framework.

As Artificial Intelligence and Machine Learning have become more ubiquitous in many parts of the society and economy, the need for transparency, fairness, and trust increases. Many of the state-of-the-art methods and algorithms are black boxes where the decision-making process is opaque to humans. Interpretable and Explainable Artificial Intelligence aims to address this issue by offering methods that either explain the decisions of black-box models or are inherently interpretable themselves. In this context, interpretable methods passively allow a human to make sense of the output while explanations are active descriptions aimed at getting a human to comprehend the output¹. Recent years have seen a focus on interpretability and explainability with a diverse set of methods being proposed to tackle different aspects of the field^{2–5}.

Prototype-based classifiers are a category of inherently interpretable methods that use prototypical examples to make their decisions. It is assumed that as long as the prototypes themselves are understandable by a human, the decision itself is interpretable⁶. Prototype-based classifiers are not new inventions. Many existed long before the need for interpretability became so urgent^{7–11}. In recent years, newer methods have been proposed that combine the power and expressability of neural networks with the decision-making process of a prototype based classifier to create prototypical neural nets^{12,13}, reaching results competitive with the state of the art while being inherently interpretable in the process.

A newer subcategory of prototype-based classifiers is part-prototype networks. These networks, usually operating in the domain of image classification use regions of a query sample, as opposed to the entire query image, to make their decisions. ProtoPNet⁴ is the first of such methods that offered fine-grained explanations for image classification while offering state-of-the-art accuracy. Figure 1 shows an example of how a part-prototype method makes its decisions.

The explanations given by part-prototype based methods can be very different from each other. Even when the general layout of the explanation is similar, the part-prototypes themselves can be vastly different. It is unusual to assume that they offer the same level of interpretability. Therefore, the evaluation of their interpretability is necessary.

Evaluation of interpretable or explainable AI methods has been a topic of research for some time. Narayanan et al.¹⁴ studied explanations in the form of decision sets to find the factors that enable humans to understand such explanations in machine learning methods. Vilone et al.¹⁵ included human-centered approaches to their study of evaluation of notions of explainability. Mohseni et al.¹⁶ offered a human-centric approach for evaluating the explanations given by methods like Grad-CAM³ and LIME². However, evaluation of part-prototype based methods requires approaches that utilize the unique architecture of these methods.

¹School of Computer Science, Carleton University, Ottawa, ON, Canada. ²Department of Computer Engineering, Sharif University of Technology, Tehran, Iran. ✉email: omid.davoudi@carleton.ca

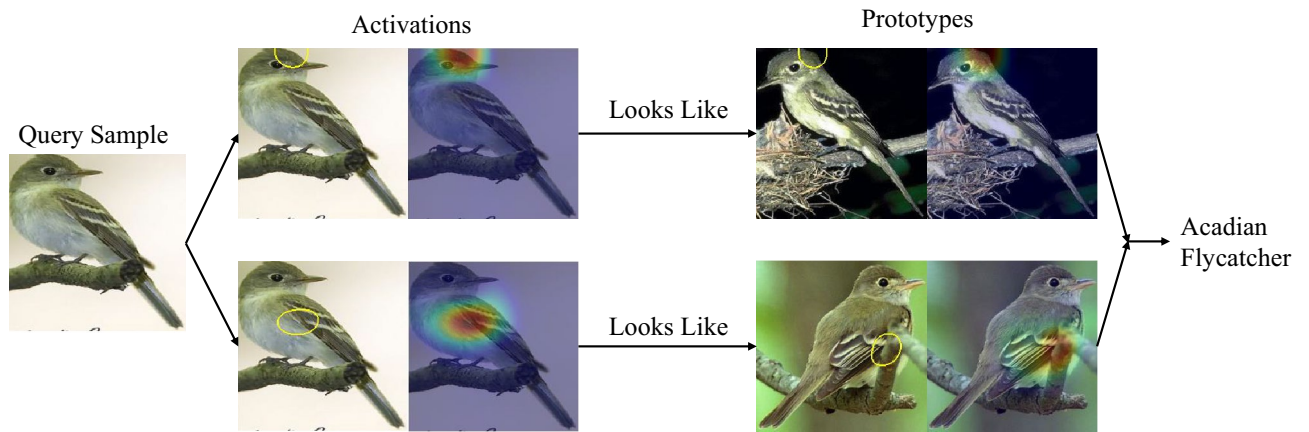


Figure 1. Example of the decision-making process of a part-prototype method.

As a result, some methods were developed specifically to evaluate interpretability in part-prototype based networks. While many of these methods evaluate the classification performance of their models and compare them to the state of the art, few analyze the interpretability of their methods. Most of the analysis in this regard seems to be focused on automatic metrics for assessing interpretability¹⁷. Such automatic metrics, while useful, are not a replacement for human evaluation of interpretability. Others have worked on human-assisted debugging¹⁸ but have not extended that to a full evaluation of method interpretability.

Kim et al. offered a method for evaluating visual concepts by humans and even performed experiments on ProtoPNet and ProtoTree¹⁹, but their evaluation suffers from a number of issues. The scale of the experiments in Kim et al. is small, with only two part-prototype methods evaluated using only a single dataset. The experimental design of that work also relies on fine-grained ratings by human annotators. This type of design can be an unreliable way of measuring human opinion when there is no consensus on what each option means²⁰. It used the class label to measure the quality of the prototypes in the CUB dataset even though there was no indication that the human users were familiar with the minutiae of the distinctions between 200 classes of birds. Lastly, it used the default rectangular representation of prototypes from ProtoPNet and ProtoTree. These representations are prone to being overly broad and misleading to the human user compared to the actual activation heatmap. As a result, we propose a human-centric analysis consisting of a set of experiments to assess the interpretability of part-prototype methods.

Goals

The interpretability of a part-prototype system is not a well-defined concept. In this work, we focus on three properties that such systems should have in order to be interpretable.

- Interpretability of the prototype itself: The concept a prototype is referring to should be recognizable and understandable to a human. Figure 2a shows an example of a prototype that is not interpretable because it points to an irrelevant background region. Machine learning methods and neural networks, in particular, can make correct decisions based on feature combinations in the data that a human might not understand. In addition, the presentation of such features is very important. A prototype might refer to a very unusual concept but its presentation might lead a human to wrongfully believe that they understand the reasoning behind a decision.

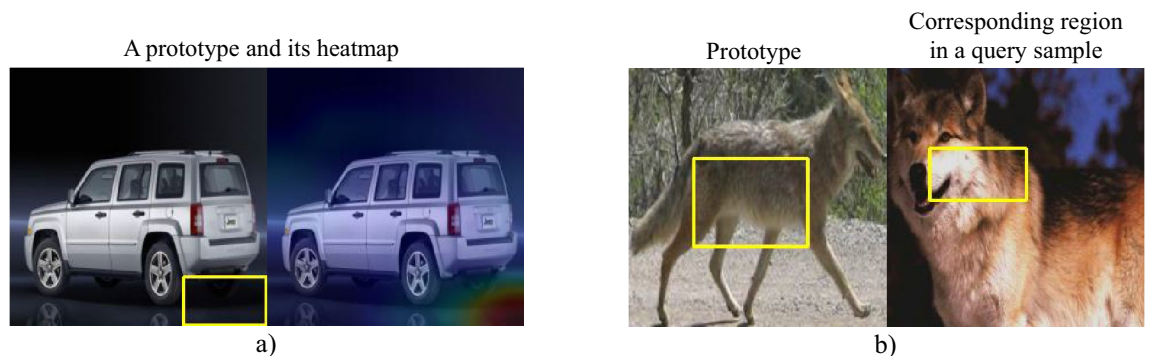


Figure 2. Examples of interpretability problems with prototypes. **(a)** The prototype itself is not interpretable because it is pointing to an irrelevant background region. **(b)** lack of similarity between a prototype and the corresponding region in the query sample.

- The similarity of a prototype to the corresponding region in the query sample: Even if the prototype itself is easily understood by a human, its activation on the query sample might not show the same concept as the prototype. Figure 2b shows an example of this problem. This is important because it shows that the structural similarity in the embedding space that the prototypes reside in is not compatible with human understanding of similarity. This is a problem that has been reported in previous literature²¹.
- The interpretability of the decision-making process itself is also an important aspect of prototype-based methods. Even if the prototypes and their similarity to the activated patches of the query sample are understood by humans, the final decision might not be. For example, a model might select and use unrelated prototypes to correctly classify a sample.

The main novelty of this work is a more robust framework for evaluating the interpretability of part-prototype-based networks using human annotators. Some previous methods have tried to do such evaluations based on automatic metrics¹⁷, and some other works have worked on human-based evaluation of interpretability for other types of explainable AI methods^{22,23}. The closest work is HIVE¹⁹ which suffers from a number of issues that are addressed in our approach. More on this will follow in the next section.

Another novelty of this work is the proposal of three actionable metrics and experiments for evaluating the interpretability of part-prototype-based classifiers. We believe that if a model fails these tests, it would not be a good interpretable model. These can assist future researchers in providing evidence rather than just making assumptions about the interpretability of their approaches.

Finally, our extensive set of experiments using Amazon Mechanical Turk includes comparisons of six related methods on three datasets. To the best of our knowledge, this is the most comprehensive work on evaluating the interpretability of such methods in a unified framework.

Background Information

HIVE

HIVE¹⁹ (Human Interpretability of Visual Explanations) is a framework for human evaluation of interpretability in AI methods that use visual explanations. HIVE uses questionnaires to ask humans to rate multiple aspects of the interpretability of AI methods. These include similarity, agreement, and distinction tests.

There are, however, some problems with the question design of the HIVE framework that makes it flawed for the purpose of assessing interpretability for prototype-based methods. In particular, they ask the human participants to rate multiple criteria in a fine-grained manner. This practice has been shown to be unreliable because human opinion can differ greatly when asked to quantify vague concepts²⁰. HIVE also assumes that every bit of explanation is interpretable by itself, even though there are examples showing the opposite (see Fig. 2a). Finally, HIVE does not follow the same process that is used by the AI model to classify a query sample. It splits the prototypes into class categories when showing them to human participants and hides the activation values. This is completely opposite of the way many of the part-prototype methods operate. Our proposed method does not suffer from these deficiencies.

ProtoPNet

ProtoPNet⁴ is a neural-network-based model that utilizes the similarity between parts of an input to a number of part-prototypes to classify images. The image first goes through a backbone network, usually a pre-trained image network such as ResNet consisting of multiple convolutional layers, and is turned into a spatially related matrix of embeddings. These embeddings are usually obtained in the penultimate layers of the common pre-trained networks. The distance between the closest of each embedding to each prototype is then fed into a fully connected output layer with softmax to determine the final class.

The resulting network is end-to-end trainable via gradient descent. A weighted sum of multiple loss functions is used to train the model. These losses include Cross-entropy loss for the final classification loss, as well as a Clustering loss designed to make sure all query samples have at least one prototype that is close to some part of them. There is also a Separation loss designed to make sure samples are far from prototypes of other classes.

SPARROW

SPARROW²⁴ is a method based on ProtoPNet that aims to make sure that the prototypes are semantically coherent. In principle, it aims to enforce sparsity, narrowness of scope and uniqueness to the prototypes. It does so by adding two additional components to the loss equation. The first one tries to decrease angular similarity between prototypes themselves to make sure they are unique, and the second one tries to make sure the prototypes themselves are as close as possible to at the query samples in the embedding space of the model in the hope of making the scope of the prototypes narrower.

Deformable ProtoPNet

Deformable ProtoPNet²⁵ is another method that aims to overcome some of the issues of the original. In this case, the rigidity of the prototypes in ProtoPNet. The prototypes in ProtoPNet are vectors with a spatially rigid semantic meaning in the pixel space. Deformable ProtoPNet aims to remedy that by creating spatially deformable prototypes where the position of features within each prototype can be different while still encoding the same concept.

ProtoTree

By leveraging the interpretability of the decision trees, a new prototype-based model was introduced called the Neural Prototype Tree (ProtoTree)²⁶. A ProtoTree consists of a Convolutional Neural Network followed by a soft binary tree structure with learnable prototypes in its decision nodes. The distance between the nearest latent patch and a prototype determines to what extent the prototype is present anywhere in an input image, which influences the routing of the image through the corresponding node. As it uses a soft decision tree, a query sample is routed through both children, each with a certain weight. After that, the model makes its prediction based on the taken paths to all leaves and their class distributions. However, the authors claim that a trained soft ProtoTree can convert into a hard ProtoTree without losing accuracy. In contrast with the ProtoPNet with a linear bag-of-prototypes, the ProtoTree enforces a sequence of steps and supports negative associations. Consequently, it can drastically reduce the number of prototypes.

TesNet

Another interpretable network architecture was designed by introducing a plug-in transparent embedding space (TesNet)²⁷ with basis concepts constructed on the Grassmann manifold. Concerning ProtoPNet, the authors mention two main disadvantages. First, the prototypes are implicitly assumed to follow a Gaussian distribution, which is improper for the complex data structure. Moreover, it cannot explicitly ensure that the learned prototypes are disentangled, an essential property for interpretable learning. The TesNet model was designed to alleviate such problems by proposing a transparent embedding space. In terms of architecture, the prototype layer in the ProtoPNet is replaced by a Transparent Subspace layer with “basis concepts”. This layer covers as many subspaces as the number of classes, and each subspace is spanned by M basis concepts. The M within-class concepts are assumed orthogonal to each other, and each concept can be traced back to the high-level patches of the feature map. To this end, it defines rigorous losses. They include orthogonality of within-class concepts, subspace separation, high-level patch grouping, and concept-based classification loss.

ProtoPool

After ProtoTree and TesNet, ProtoPool²⁸ was introduced. The model is said to have mechanisms that reduce the number of prototypes and obtain higher interpretability and easier training. The first is achieved through prototype sharing, the second comes from defining a novel similarity function, and the third results from employing a soft assignment based on prototype distributions to optimally use prototypes from the pool. The prototypes must focus on a salient visual feature, not the background. For this purpose, instead of maximizing the global activation, they widen the gap between the maximal and average similarity between the image activation map and prototypes. The other two advantages come from having a pool of prototypes and learning distributions of them per class. Therefore, the architecture is very similar to ProtoPNet, with some changes in the prototype pool layer, which contribute to prototype distribution and applying the new focal similarity function.

ACE

ACE²⁹ is an unsupervised method for extracting visual concepts from a dataset. These concepts are small regions of the images that seem to be important identifiers of the subjects of the pictures. To find these concepts, first, the images are segmented into small pieces, and then each piece is fed to a state-of-the-art classifier network. The output of the bottleneck layer of this network is then used as a representation of the concept present in the segment. Clustering methods are then utilized to find important clusters within these representations which are then picked as important concepts for the dataset.

Methodology

The interpretability of a prototype-based machine learning method as a whole is too broad to be effectively evaluated. As a result, our experiments were designed to measure the human opinion on the individual properties required for an interpretable method. This focus on single aspects allowed us to gain fine-grained insight into how much human users understand and agree with the explanations of different methods.

Human annotators were recruited via Amazon Mechanical Turk. To ensure that workers have a good understanding of the task, they were required to first complete a qualification test. Furthermore, to ensure the quality of the responses and reduce noise, we inserted a few validation samples into our tasks and excluded workers who did not pass them.

Datasets used

We used three datasets in our experiments. The datasets used had to deal with concepts that were broadly understandable to the average person. Note that it is not necessary for a human to be able to distinguish between the classes within a dataset (for example, distinguishing between different species of birds). The only requirement is for people to understand the general concept of the subjects of the image (For example, knowing what a bird is).

The datasets on which we trained our models on were CUB-200³⁰, Stanford Cars³¹, and a subset of the ImageNet dataset³² consisting of 7 classes of different felines and canines. The details for our ImageNet subset are included in the supplementary materials. All datasets were augmented in the same process described in the work by Chen et al.⁴. A concise summary of the datasets can be found in the supplements.

CUB-200-2011

CUB-200-2011³⁰ is a dataset of birds consisting of 11988 total images that are split into 200 classes. This is a popular dataset in part-prototype network research. It also depicts subjects that the average human can understand.

For our purposes, we split the dataset into two halves with 6194 images for training and 5994 images for testing. The reason for such a drastic split is that the augmentation method we used (which is the same as the one used for ProtoPNet and many of the other derived methods) multiplies the number of training images by a factor of 30.

Stanford cars dataset

The Stanford Cars Dataset³¹ is an image dataset consisting of 196 different models of cars. They are split into 8339 training and 8237 test images. This is another dataset where the subjects are generally well understood by an average human and thus, was a good pick for our experiments. The same augmentation process was also done on this dataset except that there are no bounding box coordinates available for this dataset. As a result, the images were not cropped to the subjects.

ImageNet subset

ImageNet³² is a very diverse image dataset containing a hierarchy of classes. We chose 7 classes out of these and created a subset that we believed was understandable by the average human while also containing similar concepts and a relatively balanced number of samples per class. These classes were “Big Cats”, “Collie breeds of dogs”, “Domestic Cats”, “Foxes”, “Swiss Mountain Dogs”, “Wild Cats” and “Wolves”. It contains 16586 training and 1686 test images. The same augmentation process was also done on this dataset.

Experimental setup

We trained seven methods in total. These methods include six part-prototype based classification methods (ProtoPNet⁴, Deformable ProtoPNet²⁵, SPARROW²⁴, ProtoTree²⁶, TesNet²⁷, and ProtoPool²⁸) and one unsupervised method (ACE²⁹). The unsupervised method (ACE) is normally used to extract concepts from a dataset. Each concept is a part of a training sample and therefore is analogous to the concept of prototypes in part-prototype-based methods. As a result, we decided to include it in our experiments as a method representing unsupervised approaches to this problem and also as a baseline.

We used the public code repositories from the authors of the papers of those methods in all cases except SPARROW. In the case of SPARROW, the authors kindly agreed to give us access to their non-public codebase. We used the same hyperparameters provided by these codebases, only changing those values when needed for different datasets. In particular, we used 10 prototypes for each class when training on datasets where the original codebase did not contain hyperparameters. In the special case of ProtoPool, we used 12 prototypes in total for the ImageNet dataset.

We trained all classification methods until they reached at least 70% top-1 classification accuracy on the dataset. For Imagenet, our criterion was higher at 90% due to the lower number of classes and the fact that these backbone networks were originally trained on ImageNet in the first place. All methods used ResNet-18 as their backbone network. Something to note is that ProtoPool failed to reach the 70% accuracy level for the Cars dataset, instead only reaching a Top-1 accuracy of 49.53%. Details of the performance of these models can be found in the supplements.

The next step was to gather a suite of query samples from the dataset and generate the explanations given by the methods for classifying those samples. For CUB and Cars, we used the first image found in the directory of each separate class. For Imagenet, we used the first 15 images in order to get enough samples for later use.

One thing we noticed at this stage was that the usual heatmaps provided by many of these methods to denote the location of a prototype or its activation on a query image tended to occlude a large portion of the image itself. Some, like ProtoPNet, tended to draw a rectangle to denote the general boundary of the prototype or its activation. We found this approach to be quite misleading because there were cases where the rectangles could cover more area than the actual activation site. To address this, we opted for a more fine-grained approach where we created boundaries around the area of the activation heatmap with activation values above 70%. Figure 3 illustrates examples of complex activation sites, failures of the original ProtoPNet region selection approach, and our approach to address the problem.



Figure 3. Examples of complex prototype heatmaps and the failure of a rectangular prototype representation. While the actual heatmap (left) points to irrelevant background regions, the rectangular representation (center) misleadingly includes foreground regions. Our solution (right) properly represents the actual heatmap.

Prototype interpretability

Our first set of experiments had the goal of finding out the interpretability of the prototypes created by each method. Many part prototype methods claim that their prototypes are intrinsically interpretable. In that case, humans should be able to understand the concept that is encoded within those prototypes. If a prototype represents a large area of an image without any coherent concept understandable to humans, it is not an interpretable prototype.

In the first experiment, each person was shown a pair of prototypes (concept images in the case of ACE) on the same dataset coming from different methods. The annotators had to choose the prototype they found more interpretable out of the two. Every person had to complete 10 such pairs each time. Out of those 10, 2 were validation pairs handpicked by us to ensure answer quality. An example of such a pair and the instructions given to each person are show in Fig. 4.

Only people who had finished a qualification task consisting of 7 such pairs with perfect scores were allowed into the main study. Moreover, each 10-pair set of query samples was given to three different people to gauge overall agreement and find individuals who might have chosen randomly. Additionally, this experiment was repeated 3 times and the results were aggregated.

This experiment was a purely comparative study between different methods. This ensured that people had a frame of reference when deciding which method worked better, and this is a common practice in many studies of subjective nature. But they are not without problems. One of the more important problems with this experiment was that it could not answer whether the prototypes from these methods were interpretable or not. It could only answer whether the prototypes of a particular method were *more interpretable on average* than the ones from another.

To remedy this, we conducted another set of experiments. In the new absolute experiments the annotators were given the option to also choose “Both” and “None” if they thought that both or none of the images in the pair were interpretable. The reason we did not simply show them a single image asking for interpretability was to keep the layout of the experiments very similar. This not only simplified the logistics of creating query samples for the new experiment, but also gave us the opportunity to see whether the rankings of the methods seen in the comparative experiments remained the same when the new options were added. The assumption was that if a very similar experiment with those options added resulted in a drastic change in the rankings, it would reveal problems in our methodology that were hidden beforehand. Instructions given to each person were as follows:

- For each pair of images, pick the one with the best highlighted patch. A good patch contains relevant and identifying regions/parts of the subject, and only such regions/parts.
- If both images have good or bad patches, select “Both” or “None” respectively.
- Imagine that someone points to the marked region and says: “This is an image of a bird/car/lion/etc. because of this part”. How much would you agree with that statement?

It is important to note that new image pairs were created for this experiment. This was again to see if a new set of data changes the rankings considerably, revealing any potential hidden problems with the statistical significance of the results. A new qualification task and a validation set were also created. The rest of the setup was almost identical to the comparative experiment. Each time, the annotator was given 10 prototype image pairs and told to pick the best option for each pair, with two out of those 10 pairs coming from the validation set. The experiment was repeated 4 times and the results were aggregated.

Prototype-query similarity

Our second set of experiments has the goal of finding out whether the regions activated by a prototype contain the same concept as the prototype itself. Prototype classifiers, in general, use the closest prototype to a query in their operation. Part-prototype classifiers follow the same idea by finding the closest prototype patch to another patch in the input image. For this to be interpretable, the notion of *closeness* used by the method should match that of a human. Otherwise, the explanation itself will seem nonsensical. Figure 5 shows examples of both good and bad similarity between the query patch and the prototype.

As before, the first experiment was purely comparative. Each fine-grained task consisted of two pairs of images. Each pair consisted of a prototype and its activation on a query image. The prototypes picked were the most important prototypes when trying to classify an image. The idea was to avoid prototypes that were not instrumental in the final decision as they were more likely to not be present in the query image itself. The

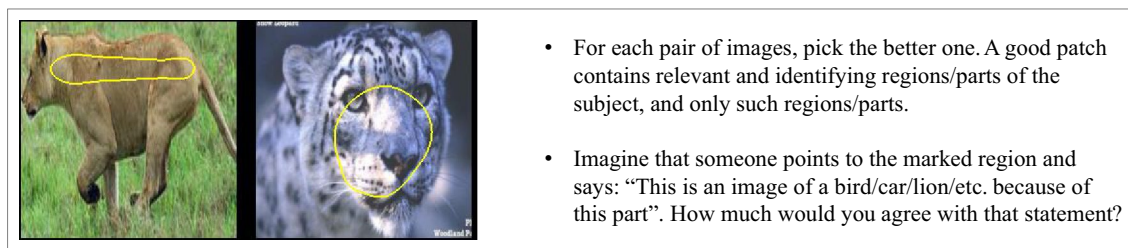


Figure 4. Examples of a prototype interpretability comparison task.

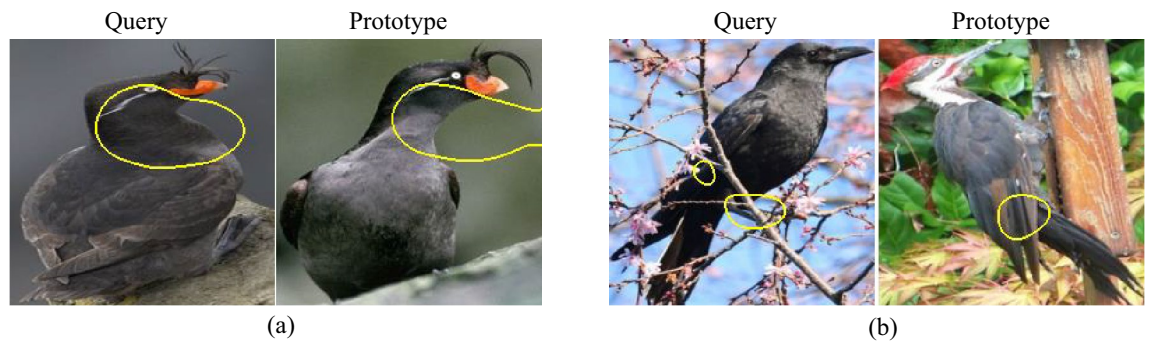


Figure 5. (a) An example of a good prototype-query patch similarity. (b) an example of a pair of prototype and query patch that are similar in the embedding space but visually very dissimilar.

annotators were then told to select the prototype/activation pair that were conceptually more similar to each other. Figure 6 shows examples of this task where annotators are expected to select pair-2 because the corresponding images are more similar than the pair-1 images.

This works well for the part-prototype classification methods as they all have an explanation generation procedure that utilizes their prototypes. On the other hand, ACE is not a classifier. It is an unsupervised concept learning method. Nevertheless, we thought that ACE by itself also had a notion of *similarity* baked into the method. It picks multiple patches from potentially different images to represent a concept. The assumption is that these patches are *similar* because they share that concept. Here, we decided to simply include two patches from the same concept in the pairs that represented ACE.

The annotators were presented with batches of 10 fine-grained tasks. Like before, 2 out of those 10 tasks contained validation data. Also similar to previous experiments, annotators had to obtain qualification by completing 7 tasks perfectly before being allowed to work on the main experiment. The experiment was repeated 3 times and the results were aggregated. The instructions shown to the annotators are shown in Fig. 6.

As with prototype interpretability experiments, we also devised a more absolute version of the similarity experiment by adding the options “Both” and “None” to the same experimental layout. The instructions for this other experiment were as follows:

- Each set contains two pairs of images. In each one, select the pair of images where the marked regions in those images are conceptually closer to each other than the other pair.
- For example if in pair 1, the marked regions are conceptually very different (for example, one is of a face and the other is of the paws of a cat), do not select that pair. On the other hand, if in pair 2, the marked regions are visually and conceptually very similar, select pair 2.
- If both pairs have similar regions, select Both. If none of the pairs has similar regions, select None.

Interpretability of the decision-making process

Our final experiment was designed to assess the interpretability of the decision-making process. In general, part-prototype methods offer a set of explanations for a decision and then use those concepts for a final classification step, which is usually a fully connected layer. We designed our framework for cases where the final decision is

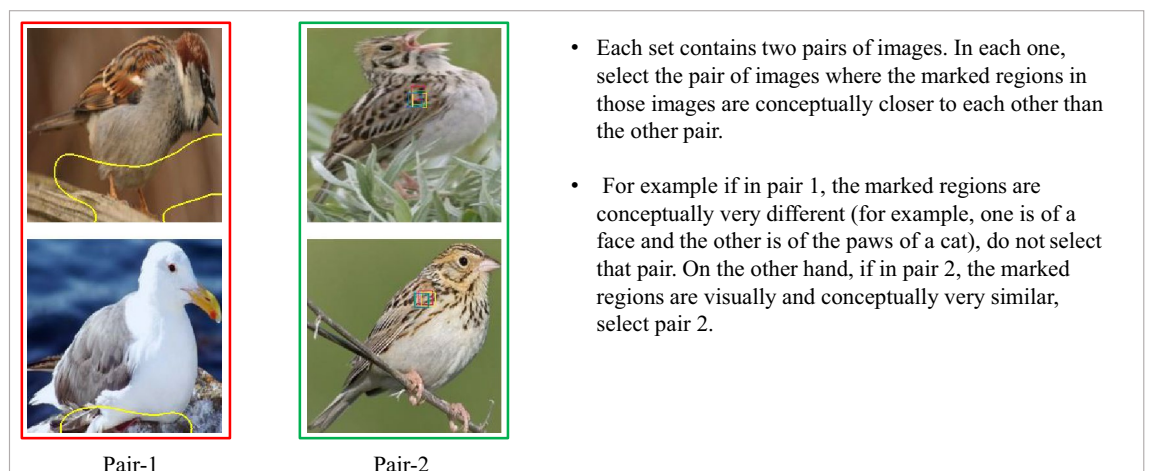


Figure 6. Two pairs of prototype and activated query patch. Annotators were told to pick the most similar pair. In this example, annotators were expected to pick Pair-2 over Pair-1.

made using such a final layer over all of the prototype activations and so we tried to stay as faithful as possible to the actual decision-making process. This meant that we had to make sure we included the activation scores and rankings of the prototypes in the experiment design.

ACE is not a classification method and there is no final decision to make. As a result, ACE, alongside ProtoTree, is not included in this experiment. In the case of ProtoTree, the method uses a decision-tree structure to offer its explanations. This is not compatible with our experiment design. Moreover, the agreement and distinction tests designed by HIVE¹⁹ are adequate for evaluating the interpretability here as they use the same tree structure as the method itself. In the end, the authors of HIVE found that the tree structure was interpretable by humans. This cannot be said for the remaining five methods.

A challenge in interpretability is to show enough information to humans so that they can understand the process without overwhelming them. For example, we cannot expect a human to look at the entire 2000 prototypes and their activations in order to understand how ProtoPNet classifies a query image from the CUB dataset. Many part-prototype models seem to have opted to use the 10 most activated prototypes for a particular query image as a good proxy for the entire model^{4,25,27}. 10 prototypes are much more manageable for a human to understand than 2000, but we have to make sure that they are also a good proxy for the model itself.

If the model, using those same 10 prototypes, is able to classify the query image and come to the same decision as before (regardless of whether that original decision was correct or not), then those 10 prototypes are a good proxy for the decision-making process of the model for that particular query image. This implies that the model selectively utilizes a limited number of prototypes from its extensive pool to categorize images.

This is unfortunately not always the case. For some query images, the model might use more than 10 prototypes to come up with its final decision. In those cases, we can argue that the interpretability of the method will suffer, as it relies on a higher number of prototypes, which, in turn, can overwhelm humans. As a result, finding out the model's accuracy with only a limited number of prototypes can be a good indicator of the interpretability of the decision-making process. We performed experiments where we only kept the top 1, 3, 5, 10, 15, 20, 30, 50, and top 100 prototypes for each method on the CUB dataset, and for each reduced model, we calculated the percentage of its decisions that agree with the decisions of the model that uses the complete prototype set. The results of these can be seen in Fig. 7.

As can be seen, not all methods are equal in this regard. Some like ProtoPNet and Deformable ProtoPNet can achieve the same decisions using relatively few prototypes. Others like ProtoPool utilize a large amount of prototypes to classify most of the query images. Only 12% of the decisions of ProtoPool remain the same when only the 10 top prototypes are kept. This observed low percentage is concerning especially as it only goes up to 39% after using 100 prototypes out of the 202 prototypes this model had for the CUB dataset.

After this step, we classified the same query samples utilized in the previous experiments using our models but only kept the ones where the decision of the model did not change after keeping only the 10 top activated prototypes. For each of these query images that passed this filter, we then gathered the top 10 prototypes and their activation regions on the query image into an explanation set. The prototypes in this explanation set were sorted from the highest activation to the lowest. As the activation values themselves might be different for each query image and might not mean anything to a human, we used a confidence system where the top activated prototype has a relative confidence level of 100% and every other prototype has its confidence level scaled by its activation value compared to the activation value of the top prototype. We then included the class label associated with that prototype in the explanation set. In the case of ProtoPool which doesn't have a single class associated with each prototype, we picked the 3 top classes associated with that prototype. These classes were obfuscated with numbers so that the role of prior human knowledge could be reduced. Examples of these explanation sets can be found in Fig. 8. We then asked our human participants on Amazon Mechanical Turk to guess the class that was picked for that query image by our classifier model. Our instructions for the AMT participants were as follows:

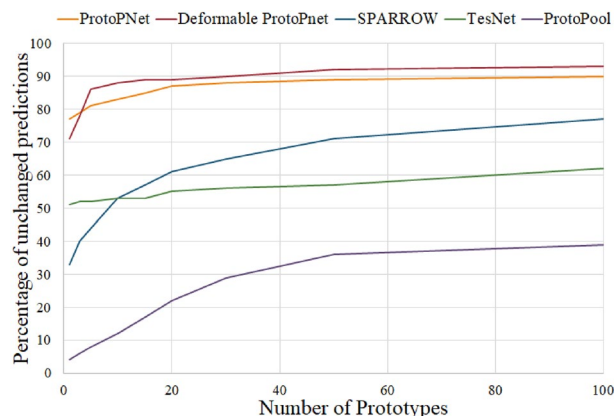


Figure 7. The degree to which a reduced model makes decisions similar to the model that uses the complete prototype set. The horizontal axis shows the number of prototypes retained in a reduced model. The vertical axis is the percentage of test data where the predictions made by the reduced model remain the same as the prediction made by the model that uses the complete prototype set on the CUB dataset.

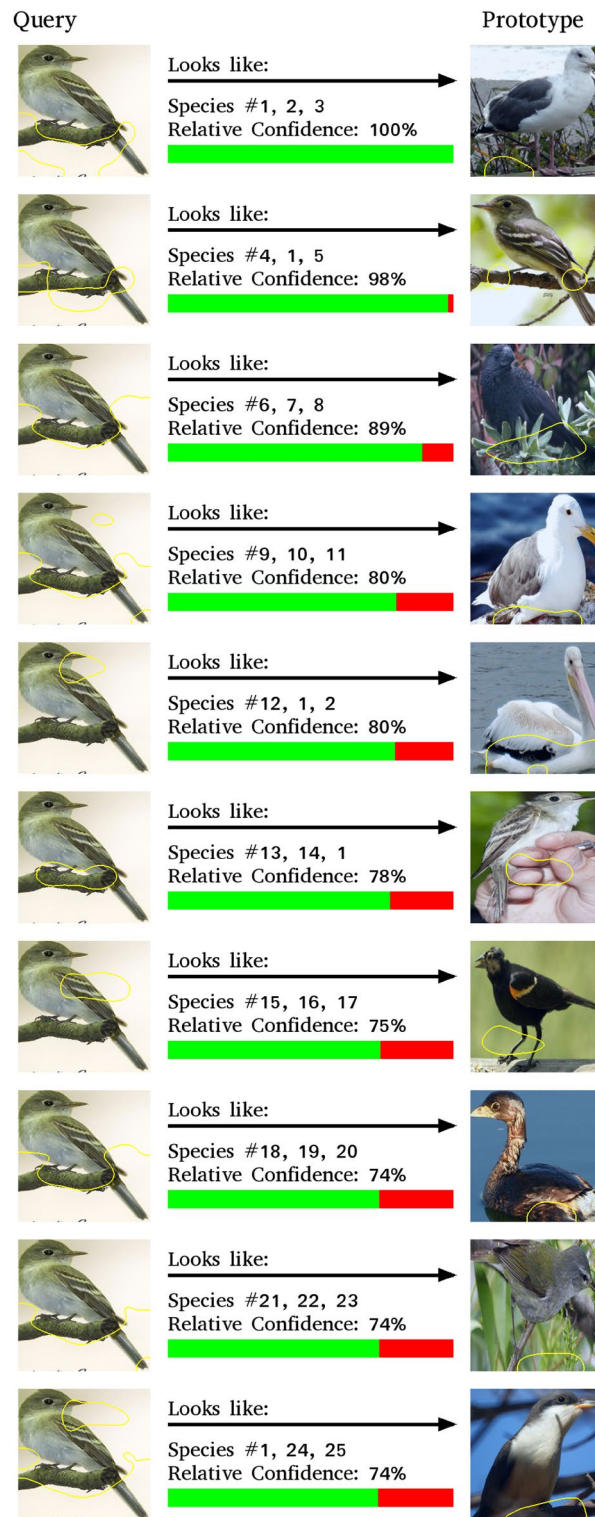


Figure 8. Example of an Explanation set used in the experiment about the interpretability of the decision-making process itself.

- Given a photo of an animal or a car (Query), an AI is predicting the species or the model of the subject of the photo using prototypes it has learned from previously seen photos. Specifically for each prototype, the AI identifies a region in the photo (Query) that looks the most similar to a region (marked with a yellow border) in the prototype and rates its relative confidence in their similarity. Each prototype is associated with one, or multiple animal species/car models. The AI will pick the species/type based on how similar and frequent the prototypes are to the query.

- For each query, we show explanations on how the AI reasons the species of the animals or the model of the car. Looking at the species/model numbers for each query-prototype pair, guess the species/model number that you would think the AI predicted for the query.
- Guessing randomly will get you low overall accuracy depending on the number of options. You will only get rewards if your performance is sufficiently higher than random guesses.
- Remember that you can zoom-in/out and pan the image using the buttons below the image.

The idea behind this experiment is similar to that of the distinction task from HIVE. In order for a human to understand the decision-making process of an AI system, they should be able to correctly predict the decision of the AI system from the same information it uses.

Results

Prototype Interpretability

In total, 1104 pairs were compared and annotated by humans for the comparative study. Table 1 shows the percentage of times each method was chosen against others. As can be seen, prototypes selected by TesNet were chosen 73.26% of the time when compared to prototypes selected by other methods. In the second place was ProtoTree with 63.64% preference. It is also interesting to note the relative weakness of SPARROW and Deformable ProtoPNet.

Table 2 shows the results of the same experiments but on a method-by-method basis. Each element of the table shows the results of the pairwise comparison between the two methods, with the method shown in the first column being the one selected out of the two. For example, the intersection between the row “ACE” and the column “ProtoPool” shows the percentage of times (42.86%) ACE was chosen against ProtoPool in the experiments. While there is more variation when looked at with this much detail, the general trends hold true.

For the absolute experiments, 1304 pairs were compared and annotated by humans in total. Table 3 shows the percentage of times each method’s prototypes were chosen in general. As the options for “Both” and “None” also existed in this experiment, the results here show the percentage of prototypes from each method that the annotators found acceptable. While the overall rankings seem to have mostly remained the same, the higher percentages overall show that while some methods might lose in the comparative study, their prototypes are still interpretable, even if a bit less so than the more successful methods. These results also show that 95.18% of the prototypes created by TesNet and 89.56% of the ones by ProtoTree are seen as interpretable. These results match our own observations, where it seems that TesNet and ProtoTree offer consistently good prototypes in all of the datasets.

Like before, Table 4 shows the results of the same experiments but on a method-by-method basis. Each row shows the percentage of times a method (shown on the first column) was selected either solely or as part of the “Both” option when paired with another method shown on the top row. Note that this time, the sum of mirrored

	ProtoPNet	Deform. PPNet	SPARROW	ACE	ProtoTree	TesNet	ProtoPool
Total Percentage Selected	51.65%	38.04%	38.69%	37.64%	63.64%	73.26%	47.41%

Table 1. Aggregated results for the comparative prototype interpretability experiment. Significant values are in [bold].

	ProtoPNet	Deform. PPNet	SPARROW	ACE	ProtoTree	TesNet	ProtoPool
ProtoPNet	–	68.93%	55.00%	66.07%	32.97%	30.48%	53.68%
Deform. PPNet	31.07%	–	53.57%	47.75%	30.28%	11.58%	48.36%
SPARROW	45.00%	46.43%	–	47.87%	37.00%	18.25%	41.73%
ACE	33.93%	52.25%	52.13%	–	21.24%	24.44%	42.86%
ProtoTree	67.03%	69.72%	63.00%	78.76%	–	43.41%	63.54%
TesNet	69.52%	88.42%	81.75%	75.56%	56.59%	–	71.26%
ProtoPool	46.32%	51.64%	58.27%	57.14%	36.46%	28.74%	–

Table 2. Results for the comparative prototype interpretability experiment.

	ProtoPNet	Deform. PPNet	SPARROW	ACE	ProtoTree	TesNet	ProtoPool
Total percentage selected	65.50%	60.26%	51.79%	53.25%	89.56%	95.18%	61.11%

Table 3. Aggregated results for the absolute prototype interpretability experiment. Significant values are in [bold].

	ProtoPNet	Deform. PPNet	SPARROW	ACE	ProtoTree	TesNet	ProtoPool
ProtoPNet	–	55.36%	67.74%	66.67%	62.50%	74.58%	65.98%
Deform. PPNet	53.57%	–	74.71%	37.04%	64.86%	57.41%	63.33%
SPARROW	38.71%	49.43%	–	37.50%	50.65%	64.13%	64.15%
ACE	46.3%	74.07%	45.83%	–	53.23%	55.81%	42.55%
ProtoTree	73.61%	89.19%	92.21%	96.77%	–	92.00%	97.92%
TesNet	89.83%	94.44%	95.65%	95.35%	98.00%	–	98.18%
ProtoPool	55.67%	50.00%	56.60%	68.09%	75.00%	69.09%	–

Table 4. Results for the absolute prototype interpretability experiment.

cells does not necessarily add up to 100% due to the existence of “Both” and “None” options. The results show a fine-grained consistency in prototype interpretability rates across all methods.

Prototype-query Similarity

In total, 1156 pairs were compared and annotated by humans for the comparative study of prototype-query similarity. Table 5 shows the percentage each method was chosen against others. As can be seen, TesNet were chosen 71.88% of times when compared to other methods. At the second place was ProtoPNet with 64.47% preference. This shows that apart from TesNet, the other methods show comparatively worse prototype-query similarity compared to ProtoPNet itself. This has an interesting implication for the interpretability of these methods compared to the baseline.

Table 6 shows the results of the same experiments but on a method-by-method basis. Each element of the table shows the results of the pairwise comparison between the two methods, with the method shown in the first column being the one selected out of the two. For example, the intersection between the row “ACE” and the column “ProtoPool” shows the percentage of times (24.17%) ACE was chosen against ProtoPool in the experiments. While there is more variation when looked at with this much detail, the general trends hold true.

For the absolute experiments, 1984 pairs were compared and annotated by humans in total. Table 7 shows the percentage of times each method’s prototypes and query sample activations were seen to be close conceptually. While the overall rankings seem to have mostly remained the same, the lower overall percentages show that none of these methods is doing well in this regard. The best results, TesNet at 42.98% and ProtoPNet at 40.41% are still quite low. This is in agreement with the results previously reported in^{19,21}.

	ProtoPNet	Deform. PPNet	Sparrow	ACE	ProtoTree	TesNet	ProtoPool
Total Percentage Selected	64.47%	46.12%	43.89%	22.63%	51.93%	71.88%	48.25%

Table 5. Aggregated results for the comparative prototype-query similarity experiment. Significant values are in [bold].

	ProtoPNet	Deform. PPNet	Sparrow	ACE	ProtoTree	TesNet	ProtoPool
ProtoPNet	–	60.17%	67.03%	87.30%	56.80%	50.38%	66.12%
Deform. PPNet	39.83%	–	48.48%	69.70%	46.00%	24.79%	51.61%
SPARROW	32.97%	51.52%	–	66.34%	41.67%	19.77%	46.88%
ACE	12.70%	30.30%	33.66%	–	23.74%	13.00%	24.17%
ProtoTree	43.20%	54.00%	58.33%	76.26%	–	33.04%	40.23%
TesNet	49.62%	75.21%	80.23%	87.00%	66.96%	–	79.82%
ProtoPool	33.88%	48.39%	53.12%	75.83%	59.77%	20.18%	–

Table 6. Results for the comparative prototype-query similarity experiment.

	ProtoPNet	Deform. PPNet	SPARROW	ACE	ProtoTree	TesNet	ProtoPool
Total percentage selected	40.41%	35.03%	36.74%	30.56%	34.97%	42.98%	31.71%

Table 7. Aggregated results for the absolute prototype-query similarity experiment. Significant values are in [bold].

Finally, Table 8 shows the results of the same experiments but on a method-by-method basis. Each row shows the percentage of times a method (shown on the first column) was selected either solely or as part of the “Both” option when paired with another method shown on the top row. Note that the sum of mirrored cells does not necessarily add up to 100% due to the existence of “Both” and “None” options. The consistently low similarity rates across different methods point to a general failure of all models in achieving a good similarity between the activated regions of the query samples and the prototypes.

Interpretability of the Decision-Making Process

In total, 607 explanation sets were shown to and annotated by the participants. Table 9 shows the percentage of query samples in which humans correctly predicted the predictions of the models. Apart from ProtoPool, human participants seem to generally be able to understand the underlying principles behind the final decision-making process of the rest of the methods. Prototypes in ProtoPool can belong to multiple classes and this causes severe issues with understanding how the method makes its decisions. The stark contrast between ProtoPool and the others shows that this evaluation criteria is a useful way of determining problems in the interpretability of Prototype-based methods.

The effects of the number of prototypes

During preparation for our experiments, we noticed an unexpected phenomenon for prototype similarity. In some dataset/method combinations, it was possible to achieve high classification accuracies of more than 80% with only one prototype per class. For example, TesNet was able to achieve about 88% accuracy when trained on the ImageNet subset using only 7 prototypes. In this instance, the similarity of a prototype with its activation was very low. This sheds light on some of the limitations of these methods.

Part-prototype-based classifiers are, in the end, discriminative models. This is in contrast to how they are perceived from an interpretability standpoint. When humans look at the explanations given for the decision, they see the similarity between the prototypes and the activated regions. However such similarity is only useful to the model itself as long as it can utilize it to discriminate effectively between the classes. As long as a query sample is closer to the prototype of the correct class than all other classes, the model makes the correct decision, regardless of the actual similarity between the query sample and the prototype.

When designing a model, setting the number of prototypes to a very small values favors interpretability. This is because it might be possible for a human to keep the entire prototype portfolio in mind when they try to understand the decision taken by the model. In other words, it may allow the human to gain a global understanding of the model’s decision as opposed to a local understanding based on only the top activated prototypes. This global interpretability comes at the expense of prototype similarity with the activation area of the query sample.

Going for higher numbers of prototypes will quickly diminish the capacity for this type of global interpretability though. As an example, keeping up with 40 prototypes instead of 20 can be quite challenging. After going over the threshold, it does not matter how much over the limit we are. Keeping a complete, global view of 100 prototypes per decision is as unfeasible for a human as that of 1000 prototypes. On the other hand, having 1000 prototypes gives the model a much better capacity to create distinct prototype patches which in turn, increases the likelihood of finding a genuinely similar prototype to a query sample. This might lead to the belief that picking the most number of prototypes possible based on the computational or memory resources available might be the best choice. However, the issue of prototype duplication³⁵ will cause diminishing returns for this strategy. To achieve the best balance, it might be best to monitor the overall number of unique prototypes while using the evaluation metrics proposed in this paper to find a sweet spot between interpretability and resource usage.

	ProtoPNet	Deform. PPNet	SPARROW	ACE	ProtoTree	TesNet	ProtoPool
ProtoPNet	–	39.56%	35.29%	31.18%	46.23%	49.02%	37.66%
Deform. PPNet	35.16%	–	28.07%	28.92%	33.98%	52.53%	32.20%
SPARROW	30.88%	47.37%	–	27.62%	30.67%	35.35%	43.10%
ACE	29.03%	31.33%	33.33%	–	31.71%	25.93%	31.19%
ProtoTree	36.79%	32.04%	34.67%	29.27%	–	40.37%	35.14%
TesNet	44.12%	46.46%	31.31%	35.80%	53.21%	–	45.00%
ProtoPool	33.77%	27.97%	32.76%	26.61%	43.24%	30.00%	–

Table 8. Results for the absolute prototype-query similarity experiment.

	ProtoPNet	Deform. PPNet	SPARROW	TesNet	ProtoPool
Percentage of correct predictions	96.01%	89.92%	85.40%	92.86%	48.77%

Table 9. Results for the interpretability of the decision-making process. Numbers are the percentage of model decisions that human participants could correctly predict based on the explanation given by the model.

Discussions

In a substantial portion of the previous work, it was implicitly assumed that part-prototype networks are interpretable by nature. We have devised three human-centric evaluation schemes to assess that assumption. One for evaluating the interpretability of the prototypes themselves, one for evaluating the similarity between prototypes and the activated regions of the query sample, and one for evaluating the interpretability of the decision-making process itself. Our experiments show that these schemes are able to differentiate between various methods in terms of interpretability while not suffering from the problems of the previous works. Moreover, we applied this scheme to seven related methods on three datasets. The results shed light on the interpretability of these methods from a human perspective.

The results show that not all part-prototype methods are equal when it comes to human interpretability. In some cases, there are severe issues that hurt the interpretability of these methods. Chief among them is the dissimilarity between the prototype and the activation region of the query. This problem existed for all of the models tested to some degree. ProtoPNet itself and TesNet performed the best among the tested methods. Deformable ProtoPNet and SPARROW have smaller activation patches which decrease the likelihood of the patches having incidental similarity in the peripheries, decreasing the overall similarity rates. ProtoTree creates a series of discriminative decisions for its operation which means that the similarity itself is less important as opposed to the discriminative ability of the prototype. ProtoPool suffers from having less prototypes compared to the other methods and even sharing those between the different classes. This decreases the chance for a similar prototype to be found for each query sample.

In addition to the problems with the prototype and query sample similarity, the prototypes of some models were not sufficiently interpretable. TesNet and ProtoTree performed the best among the methods. TesNet's emphasis on the transparent embedding space results in a very good prototype selection. ProtoTree needs to classify images using only a limited sequence of decisions based on a limited set of prototypes. This naturally forces it to pick important discriminative concepts as its prototypes, otherwise its performance would suffer significantly. ProtoPool tends to pick overly broad prototypes as it uses a smaller, shared, set of prototypes as opposed to a larger, class-specific set of the other methods. This reduces its prototype interpretability. Deformable ProtoPNet has the opposite problem. Its prototype representation consists of very small patches which can contain very little information as a whole. In the case of SPARROW, its L^{PSD} loss term will force the prototypes to converge on frequently seen patches of inputs. This usually causes the prototypes to pick frequently repeated background sections instead of the subject itself, degrading prototype interpretability significantly.

Finally, ProtoPool has a noticeable problem with the interpretability of the decision-making process due to assigning multiple classes to each prototype and requiring huge numbers of prototypes to make their final decisions. The decision-making process of some other methods were easy to understand using the top 10 activated prototypes, but they still needed more than 10 prototypes to make their final decision.

Still, in some other cases, the methods performed relatively well. Apart from ProtoPool, all methods performed quite well in the interpretability of the decision-making process test. ProtoPNet and Deformable ProtoPNet were also able to make the majority of their predictions using the top 10 activated prototypes. Prototype interpretability was also particularly high for some of the methods involved. Considering these results and also the observations on low prototype count models, our suggestion is to look at the top prototype sets for the final decision rather than only individual prototypes to get a better understanding of a model. Even if the most activated prototype has low interpretability or is dissimilar to its activation, the other top prototypes used in the decision could better explain the final decision. It is important to note that the opposite could also happen: the top prototype being interpretable but the others showing a flawed decision-making process.

The interpretability of many machine learning methods used in production is very important. However, it is more important to understand the limitations and peculiarities of the models, methods, and tools used to address the interpretability problem. In particular, unified frameworks with emphasis on human-centric assessments are necessary if we want to truly evaluate interpretability methods. We think that more studies like this should be done on different areas of AI interpretability to further our understanding of such models.

Data availability

Anonymized data of the experiments can be found in the following repository: <https://github.com/omiddavoodi/part-prototype-interpretability-data>.

Received: 12 October 2023; Accepted: 12 December 2023

Published online: 28 December 2023

References

1. Broniatowski, D. A. *et al.* *Psychological Foundations of Explainability and Interpretability in Artificial Intelligence* (Technology Representatives Inc., NIST, 2021).
2. Ribeiro, M. T., Singh, S. & Guestrin, C. “why should i trust you?” explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1135–1144 (2016).
3. Selvaraju, R. R. *et al.* Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE International Conference on Computer Vision*, 618–626 (2017).
4. Chen, C. *et al.* This looks like that: deep learning for interpretable image recognition. *Adv. Neural Inf. Process. Syst.* **32** (2019).
5. La Gatta, V., Moscato, V., Postiglione, M. & Sperli, G. PASTLE: Pivot-aided space transformation for local explanations. *Pattern Recogn. Lett.* **149**, 67–74. <https://doi.org/10.1016/j.patrec.2021.05.018> (2021).
6. Molnar, C. *Interpretable machine learning* (Lulu.com, 2020).
7. Bezdek, J. C. & Castelaz, P. F. Prototype classification and feature selection with fuzzy sets. *IEEE Trans. Syst. Man Cybern.* **7**, 87–92 (1977).

8. Kohonen, T. Improved versions of learning vector quantization. In *1990 IJCNN International Joint Conference on Neural Networks*, 545–550 (IEEE, 1990).
9. Kuncheva, L. I. & Bezdek, J. C. Nearest prototype classification: Clustering, genetic algorithms, or random search?. *IEEE Trans. Syst. Man. Cybern. Part C (Appl. Rev.)* **28**, 160–164 (1998).
10. Seo, S., Bode, M. & Obermayer, K. Soft nearest prototype classification. *IEEE Trans. Neural Netw.* **14**, 390–398 (2003).
11. Graf, A. B., Bousquet, O., Rätsch, G. & Schölkopf, B. Prototype classification: Insights from machine learning. *Neural Comput.* **21**, 272–300 (2009).
12. Li, O., Liu, H., Chen, C. & Rudin, C. Deep learning for case-based reasoning through prototypes: A neural network that explains its predictions. In *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 32 (2018).
13. Davoudi, S. O. & Komeili, M. Toward faithful case-based reasoning through learning prototypes in a nearest neighbor-friendly space. In *International Conference on Learning Representations* (2021).
14. Narayanan, M. *et al.* How do humans understand explanations from machine learning systems? An evaluation of the human-interpretability of explanation. arXiv preprint [arXiv:1802.00682](https://arxiv.org/abs/1802.00682) (2018).
15. Vilone, G. & Longo, L. Notions of explainability and evaluation approaches for explainable artificial intelligence. *Inf. Fusion* **76**, 89–106 (2021).
16. Mohseni, S., Block, J. E. & Ragan, E. Quantitative evaluation of machine learning explanations: A human-grounded benchmark. In *26th International Conference on Intelligent User Interfaces*, 22–31 (2021).
17. Huang, Q. *et al.* Evaluation and improvement of interpretability for self-explainable part-prototype networks (2023). [arXiv:2212.05946](https://arxiv.org/abs/2212.05946).
18. Bontempelli, A., Teso, S., Tentori, K., Giunchiglia, F. & Passerini, A. Concept-level debugging of part-prototype networks. arXiv preprint [arXiv:2205.15769](https://arxiv.org/abs/2205.15769) (2022).
19. Kim, S. S., Meister, N., Ramaswamy, V. V., Fong, R. & Russakovsky, O. Hive: Evaluating the human interpretability of visual explanations. In *European Conference on Computer Vision*, 280–298 (Springer, 2022).
20. Krosnick, J. A. *Questionnaire Design* 439–455 (Springer International Publishing, 2018).
21. Hoffmann, A., Fanconi, C., Rade, R. & Kohler, J. This looks like that... does it? shortcomings of latent space prototype interpretability in deep networks. arXiv preprint [arXiv:2105.02968](https://arxiv.org/abs/2105.02968) (2021).
22. Lage, I., Ross, A., Gershman, S. J., Kim, B. & Doshi-Velez, F. Human-in-the-loop interpretability prior. *Adv. Neural Inf. Process. Syst.* **31** (2018).
23. Colin, J., Fel, T., Cadène, R. & Serre, T. What i cannot predict, i do not understand: A human-centered evaluation framework for explainability methods. *Adv. Neural Inf. Process. Syst.* **35**, 2832–2845 (2022).
24. Kraft, S. *et al.* Sparrow: Semantically coherent prototypes for image classification. In *The 32nd British Machine Vision Conference (BMVC)* (2021).
25. Donnelly, J., Barnett, A. J. & Chen, C. Deformable protopnet: An interpretable image classifier using deformable prototypes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10265–10275 (2022).
26. Nauta, M., Van Bree, R. & Seifert, C. Neural prototype trees for interpretable fine-grained image recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14933–14943 (2021).
27. Wang, J., Liu, H., Wang, X. & Jing, L. Interpretable image recognition by constructing transparent embedding space. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 895–904 (2021).
28. Rymarczyk, D. *et al.* Interpretable image classification with differentiable prototypes assignment. In *European Conference on Computer Vision*, 351–368 (Springer, 2022).
29. Ghorbani, A., Wexler, J., Zou, J. Y. & Kim, B. Towards automatic concept-based explanations. *Adv. Neural Inf. Process. Syst.* **32** (2019).
30. Wah, C., Branson, S., Welinder, P., Perona, P. & Belongie, S. Cub-200-2011 dataset. Technical Report CNS-TR-2011-001, California Institute of Technology (2011).
31. Krause, J., Stark, M., Deng, J. & Fei-Fei, L. 3D object representations for fine-grained categorization. In *2013 IEEE International Conference on Computer Vision Workshops*, 554–561. <https://doi.org/10.1109/ICCVW.2013.77> (2013).
32. Deng, J. *et al.* Imagenet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, 248–255 (IEEE, 2009).

Author contributions

O.D. Designed, developed, and implemented the experiments as well as wrote the paper, S.M. developed and implemented the experiments, and M.K. designed the experiments and wrote the paper.

Competing Interests

The authors declare no competing interests.

Additional information

Correspondence and requests for materials should be addressed to O.D.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2023