

---

# What I Cannot Predict, I Do Not Understand: A Human-Centered Evaluation Framework for Explainability Methods

---

Julien Colin<sup>1,3,5 \* †</sup>   Thomas Fel<sup>1,3,4 \*</sup>   Rémi Cadène<sup>1,2 ‡</sup>   Thomas Serre<sup>1,3</sup>

<sup>1</sup>Carney Institute for Brain Science, Brown University, USA   <sup>2</sup>Sorbonne Université, CNRS, France

<sup>3</sup>Artificial and Natural Intelligence Toulouse Institute, Université de Toulouse, France

<sup>4</sup>Innovation & Research Division, SNCF

<sup>5</sup>ELLIS Alicante, Spain

{julien\_colin, thomas\_fel, remi\_cadene}@brown.edu

## Abstract

A multitude of explainability methods has been described to try to help users better understand how modern AI systems make decisions. However, most performance metrics developed to evaluate these methods have remained largely theoretical – without much consideration for the human end-user. In particular, it is not yet clear (1) how useful current explainability methods are in real-world scenarios; and (2) whether current performance metrics accurately reflect the usefulness of explanation methods for the end user. To fill this gap, we conducted psychophysics experiments at scale ( $n = 1,150$ ) to evaluate the usefulness of representative attribution methods in three real-world scenarios. Our results demonstrate that the degree to which individual attribution methods help human participants better understand an AI system varies widely across these scenarios. This suggests the need to move beyond quantitative improvements of current attribution methods, towards the development of complementary approaches that provide qualitatively different sources of information to human end-users.

## 1 Introduction

There is now broad consensus that modern AI systems might not be safe to be deployed in the real world [1] despite their exhibiting very high levels of accuracy on held-out data because these systems have been shown to exploit dataset biases and other statistical shortcuts [2–5]. A growing body of research thus focuses on the development of explainability methods to help better interpret these systems’ predictions [6–14] to make them more trustworthy. The application of these explainability methods will find broad societal uses, like easing the debugging of self-driving vehicles [15] and helping to fulfill the “right to explanation” that European laws guarantee to its citizens [16].

The most commonly used methods in eXplainable AI (XAI) are attribution methods, but despite a plethora of different approaches [6–10, 12, 17–19], assessing the quality and reliability of these methods remains an open problem. So far the community has mostly focused on evaluating these methods using surrogate measures defined axiomatically [15, 20–23]. The two most popular approaches are: (1) using ground truth annotations [9, 10, 12, 24–27] and (2) measuring *faithfulness* using objective metrics [9, 12, 24, 25, 28, 29].

The first approach takes humans into consideration only in so far as it evaluates the alignment of an explanation with a human annotation. This metric assumes that the classifier relies on human-

---

\* Equal contribution   † Work done while the author was working at Brown University   ‡ Work done before April 2021 and joining Tesla

like features\*, which is an assumption that turns out to be erroneous [2–4], and unsurprisingly the results obtained on those benchmarks do not correlate with the ability of attribution methods to help humans in decision making tasks [30]. On the other hand, the second approach focuses purely on the relationship between the explanations and the model, i.e., they assess whether the explanations actually reflect the evidence for the prediction. They have emerged as the primary means of evaluating attribution methods, but because they take the human out-of-the-loop of the evaluation, it is not clear how well they can predict the practical usefulness of attribution methods and their possible failures.

As previously highlighted in [30], relying solely on current theoretical benchmarks to evaluate attribution methods can be a risky endeavor as they are totally detached from humans. Indeed, the seminal work of [31] has emphasized keeping in mind the end goal of XAI: to develop useful methods, i.e., methods that help the user better understand their model. Thus, they recommend systematically having recourse to human experiences. Moreover, a large body of research from the psychology literature has endorsed the argument that the functional role of explanations is to support learning and generalization [32–38] – i.e., can an explanation help identify generalizable rules that readily transfer to unseen instances? Therefore, evaluating attribution methods on their ability to help users understand general rules to predict a classifier decisions is perfectly aligned with this vision.

To that end, [31] suggest evaluating explainability methods directly through the end goals of XAI. We subscribe to this idea and, in this work, consider three main real-world scenarios that illustrate potential applications of explainability methods: (1) identifying a potential bias in a system’s decision along with its source using the classic Husky vs. Wolf dataset [6] (previous work has documented the risk associated with a biased model [39–41]); (2) identifying novel strategies discovered by a system [42, 43] for tasks that would be nearly impossible for humans [44–47], here we consider a highly complex categorization problem for non-expert humans [48, 49]; and (3) understanding failure cases [4, 39] on a subset of ImageNet [50].

The main contributions of this paper are as follows:

- We propose a novel human-centered explainability performance measure together with associated psychophysics methods to experimentally evaluate the practical usefulness of modern explainability methods in real-world scenarios.
- Large-scale psychophysics experiments ( $n = 1,150$ ) revealed that SmoothGrad [8] is the most useful attribution method amongst all tested and that none of the faithfulness performance metrics appear to predict if and when attribution methods will be of practical use to a human end-user.
- Perceptual scores derived from attribution maps, characterizing either the complexity of an explanation or the challenge associated with identifying “what” features drive the system’s decision, appear to predict failure cases of explainability methods better than faithfulness metrics.

## 2 Related work

**Evaluations based on faithfulness measures** Common approaches [9, 28] measure the faithfulness of an explanation through the change in the classification score when the most important pixels are progressively removed. The bigger the drop, the more faithful is the explanation method. To ensure that the drop in score does not come from a change in the distribution of the perturbed images, the ROAR [51] methods include an additional step whereby the image classifier is re-trained between each removal step. Because these methods do not require ground-truth annotations (i.e. object masks or bounding boxes), they are quite popular in computer vision [9, 12, 24, 25, 28, 29] and natural language processing [12, 52, 53].

Nevertheless, *faithfulness* measures have recently been criticized as they all rely on a baseline for removing important areas, a baseline that will obviously give better scores to methods relying internally on the same baseline [54]. More importantly, they do not consider humans at any time in the evaluation. As a result, it is unclear if the most faithful attribution method is practically useful to humans.

**Evaluations based on humans** A second class of approaches consists in evaluating the ability of humans to leverage explanations for different purposes [6, 10, 30, 55–62]. [6] were the first to evaluate the usefulness of explanations. Their work focused on the use case of bias detection: they

---

\*One advantage of our approach is that it is independent of the classifier, whether it relies on human-like features or it is even correct.

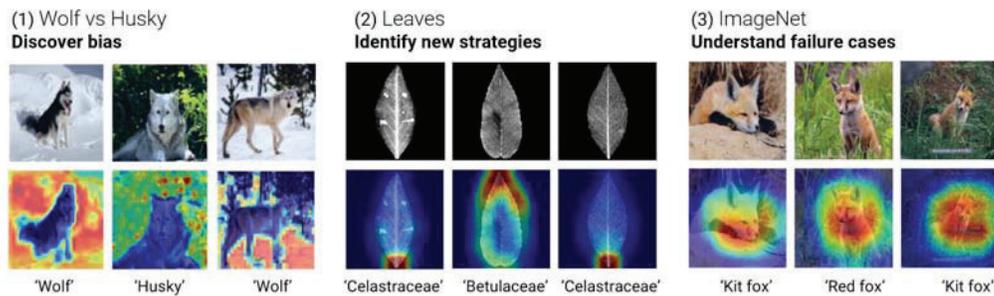


Figure 1: We study the practical usefulness of recent explainability methods in three real-world scenarios, each corresponding to different use cases for XAI. The first dataset is Husky vs. Wolf where the goal of the explanations is to help the user to identify a source of bias in a model (classification is based on the background (snow, grass) as opposed to the animal). The second dataset corresponds to a real-world leaf classification problem which is complex for non-experts. The goal of the explanations is to help the end-user identify the strategy discovered by the vision system. Finally, the third dataset is a subset of ImageNet, which consists of a collection of images where half have been misclassified by the system. The goal of the explanations here is to help the end-user understand the failure sources of a high performing model.

trained a classifier on a biased dataset of wolves and huskies and found that the model consistently used the background to classify. They asked participants if they trusted the model before and after seeing the explanation for the model's predictions, and found that explanations helped detect bias here. We use a similar dataset to reproduce those results, but our evaluation differs greatly from theirs as we do not ask if participants trust the model but instead measure directly if they understand it.

Closest to our work are [59–61, 63, 64]. [59–61] design their evaluation around the notion of *simulatability* [31, 65]. They introduce different experimental procedures to measure if humans can learn from the explanations how to copy the model prediction on unseen data. Some provide the explanations at test time [59, 60]. Similar to us but for tabular data, [61] proposes to hide explanations at test time, this forces the participants to learn the rules driving the model's decision at training time where the explanations are shown. There are two limitations to their work: (1) they provide ground-truth labels associated with input images during training, (2) the participants see the same set of images without explanations, and then with explanations, always in that order. This creates learning effects that can heavily bias their results. We differ from their work by: (1) removing ground-truth labels from our framework as they serve no purpose and can bias participants, and (2) we have different participants go through the different conditions. This removes any learning effect, and more importantly, new explainability methods can be evaluated independently and still be compared to the previously evaluated methods. A recent study [63] evaluated how AI systems may be able to assist human decisions by asking participants to identify the correct prediction out of four prediction-explanations pairs shown simultaneously. This measure reflects how well explanations help users tell apart correct from incorrect predictions. While the approach was useful to evaluate explanations in this specific scenario, it is not clear how this framework could be used to evaluate explainability methods more generally. Furthermore, when comparing different types of methods, they adapt the complexity of certain explainability methods to ease the task for participants. We argue that the complexity of explanations is an important property of explanations and that abstracting it away from the evaluation lead to unfair comparisons between methods. In contrast, we propose a more general evaluation framework that can be used for any kind of explainability method without the need to adapt them for the evaluation procedure – hence allowing for an unbiased and scalable comparison between methods. Finally, [64] proposes to evaluate if users are able to identify important features biasing the predictions of a model using a synthetic dataset. By controlling the generation process of the dataset, they have access to the ground-truth attributes biasing the classifier, and can measure the accuracy of users at identifying these features. They evaluate if concept-based or counterfactual explanations help users improve over a baseline accuracy when no explanations are provided, and find no explanation tested to be useful. While both works highlights the importance of human evaluation, they differ in: the metrics employed (identifying relevant features for the model vs. meta-prediction), the type of dataset used (synthetic vs. real-world scenarios), and the type of methods evaluated (counterfactual and concept-based methods vs. attribution methods).

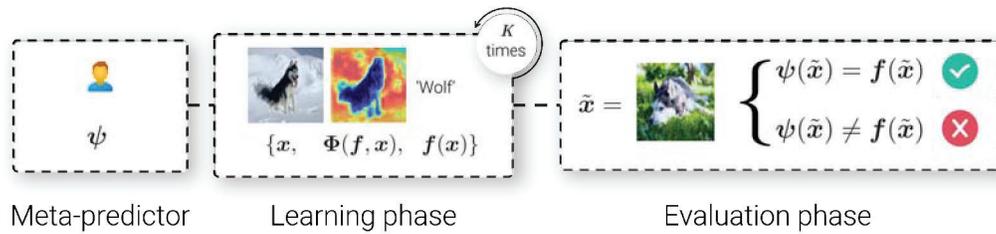


Figure 2: We describe a human-centered framework to evaluate explainability methods borrowing the concept of Meta-predictor. The framework requires a black box model  $f$  (the predictor), an explanation method  $\Phi$  and a human subject  $\psi$  which will try to predict the predictor, hence, the name Meta-predictor. The first step is the learning phase where the Meta-predictor is training using  $K$  samples  $x$ , together with the associated model predictions  $f(x)$  and explanations  $\Phi(f, x)$ . The goal of this learning phase is for the Meta-predictor to uncover the rules driving the decisions of the model from the triplets  $(x, \Phi(f, x), f(x))$ . Then, the second step is the evaluation phase where we test the Meta-predictor's ability to correctly predict the model's outputs on new samples  $\tilde{x}$  by comparing its predictions  $\psi(\tilde{x})$  to those of  $f(\tilde{x})$ . The *Utility* score of the explanation method is then computed as the relative accuracy improvement of Meta-predictor trained with vs. without explanations.

### 3 Proposed evaluation framework

Before providing a rigorous definition of interpretability, let us motivate our approach with an example: a linear classifier is often considered to be readily interpretable because its inner working is sufficiently intuitive that it can be comprehended by a human user. A user can in turn build a mental model of the classifier – predicting the classifier's output for arbitrary inputs. In essence, we suggest that the model is interpretable because the output can be predicted – i.e, we say we understand the rules used by a model, if we can use those inferred rules to correctly predict its output. This concept of predicting the classifier's output is central to our approach and we conceptualize the human user as a *Meta-predictor* of the machine learning model. This notion of Meta-predictor is also closely related to the notion of *simulatability* [24, 31, 61, 65, 66]. We will now define the term more formally.

We consider a standard supervised learning setting where  $f$  is a black-box predictor that maps an input  $x \in \mathcal{X}$  (e.g., an image) to an output  $f(x) \in \mathcal{Y}$  (e.g., a class label). One of the main goals of eXplainable AI is to yield useful rules to understand the inner-working of a model  $f$  such that it is possible to infer its behavior on unseen data points. To correctly infer those rules, the usual approach consists in studying explanations (from Attribution Map, Concept Activation Vectors, Feature Visualization, *etc.*) for several predictions. Formally,  $\Phi$  is any explanation functional which, given a predictor  $f$  and a point  $x$ , provides an information  $\Phi(f, x)$  about the prediction of the predictor. In our experiments,  $\Phi$  is an attribution method but we would like to remind that the framework is naturally adaptable to other explainability methods such as concept-based methods or feature visualization.

**The understandability-completeness trade-off** Different attribution methods will typically produce different heatmaps – potentially highlighting different image regions and/or presenting the same information in a different format. The quality of an explanation can thus be affected by two factors: *faithfulness* of the explanation (i.e., how many pixels or input dimensions deemed important effectively drive the classifier's prediction) and the understandability of the explanation for an end-user (i.e., how much of the pattern highlighted by the explanation is grasped by the user).

At one extreme, an explanation can be entirely *faithful* and provide all the information necessary to predict how a classifier will assign a class label to an arbitrary image (i.e., by giving all the parameters of the classifiers). However, such information will obviously be too complex to be understood by a user and hence it is not *understandable*. Conversely, an explanation that overly simplifies the model might offer an approximation of the rule used by the model that will be more easily grasped by the user – a more *understandable* explanation – but this approximation might ultimately mislead the user if it is not *faithful*. That is to say, just because a human agrees with the evidence pointed out by an explanation does not necessarily mean that it reflects how the model works.

Overall, this means that there is a trade-off between the amount of information provided by an explanation and its comprehensibility to humans. The most useful explanations should lie somewhere in the middle of this trade-off.

**The usefulness metric** We describe a new human-centered measure that incorporates this trade-off into a single *usefulness* measure by empirically evaluating the ability of human participants to learn to “predict the predictor”, i.e., to be an accurate Meta-predictor. Indeed, if an explanation allows users to infer precise rules for the functioning of the predictor on past data, the correct application of these same rules should allow the user to correctly anticipate the model’s decisions on future data. Scrutable but inaccurate explanations will result in an inaccurate Meta-predictor – just like accurate inscrutable ones. This Meta-predictor framework avoids current pitfalls such as confirmation bias – just because a user likes the explanation does not mean they will be a better Meta-predictor – or prediction leakage on the explanation – in simulatability experiments, as the explanation is available during the test phase, any explanation that leaks the prediction would have a perfect score, without giving us any additional information about the model. We will now formally describe the metric build using this framework.

We assume a dataset<sup>†</sup>  $\mathcal{D} = \{(\mathbf{x}_i, \mathbf{f}(\mathbf{x}_i), \Phi(\mathbf{f}, \mathbf{x}_i))\}_{i=1}^K$  used to train human participants to learn to predict a classifier’s output  $\mathbf{f}$  from  $K$  samples made of an input image  $\mathbf{x}_i$ , the associated predictions  $\mathbf{f}(\mathbf{x}_i)$  and explanations  $\Phi(\mathbf{f}, \mathbf{x}_i)$ . We denote  $\psi^{(K)}$  a human Meta-predictor after being trained on the dataset  $\mathcal{D}$  (see Fig. 2) using explanations. In addition, let  $\psi^{(0)}$  be the human Meta-predictor after participants were trained on the same dataset but without explanations to offer baseline accuracy scores. We can now define the usefulness of an explainability method  $\Phi$  after training participants on  $K$  samples through the accuracy score of the Meta-predictor normalized by the baseline Meta-predictor accuracy:

$$Utility-K = \frac{\mathbb{P}(\psi^{(K)}(\mathbf{x}) = \mathbf{f}(\mathbf{x}))}{\mathbb{P}(\psi^{(0)}(\mathbf{x}) = \mathbf{f}(\mathbf{x}))} \quad (1)$$

with  $\mathbb{P}(\cdot)$  the probability over a test set. Thus, *Utility-K* score measures the improvement in accuracy that the explanation has brought. It is important to emphasize that this *Utility* measure only depends on the classifier prediction and not on the ground-truth label as recommended by [67]. After fixing the number of training samples  $K$ , we compare the normalized accuracy of different Meta-predictors. The Meta-predictor with the highest score is then the one whose explanations were the most useful as measures compared to a no-explanation baseline.

**Utility metric** In practice, we propose to vary the number of observations  $K \in \{K_0, \dots, K_n\}$  and to report an aggregated *Utility* score by computing the area under the curve (AUC) of the *Utility-K*. The higher the AUC the better the corresponding explanation method is. Formally, given a curve represented by a set of  $n$  points  $\mathcal{C} = \{(K_0, Utility-K_0), \dots, (K_n, Utility-K_n)\}$  where  $K_{i-1} < K_i$  we define the metric as  $Utility = AUC(\mathcal{C})$ .

## 4 Experimental design

We first describe how participants were enrolled in the study, then our general experimental design (See SI for more informations).

**Participants** Behavioral data were gathered from  $n = 1,150$  participants using Amazon Mechanical Turk (AMT) ([www.mturk.com](http://www.mturk.com)). All participants provided informed consent electronically and were compensated \$1.4 for their time ( $\sim 5 - 8$  min). The protocol was approved by the University IRB and was carried out in accordance with the provisions of the World Medical Association Declaration of Helsinki. For each of the three tested datasets, we ensured that there was a sufficient number of participants after filtering out uncooperative participants ( $n = 240$  participants, 30 per condition, 8 conditions) to guarantee sufficient statistical power (See SI for details). Overall, the cost of evaluating one method using our benchmark is relatively modest (\$50 per test scenario).

**General study design** It included 3 conditions: an experimental condition where an explanation is provided to human participants during their training phase (see Fig. 2), a baseline condition where no explanation was provided to the human participants, and a control condition where a bottom-up saliency map [68] was provided as a non-informative explanation. This last control is critical, and indeed lacking from previous work [6, 61], because it provides a control for the possibility that

<sup>†</sup>We note that, in this paper, we only considered binary dataset –Class 1 vs Class 2– because having the participants classify more than 2 classes would increase their cognitive load and bring unnecessary difficulty to the task. Nonetheless, any dataset could have been used as classification problems with more than 2 classes can always be trivially reformulated as Target class vs. Other / binary classification problems, instead of Class 1 vs Class 2, without lack of generality.

Method	Husky vs. Wolf				Leaves				ImageNet			
	1	2	3	Utility	1	2	3	Utility	1	2	3	Utility
Baseline	55.7	66.2	62.9		70.1	76.8	78.6		58.8	62.2	58.8	
Control	53.3	61.0	61.4	0.95	72.0	78.0	80.2	1.02	60.7	59.2	48.5	0.94
Saliency [17]	53.9	69.6	73.3	1.06	83.2	88.7	82.4	<b>1.13</b>	61.7	60.2	58.2	1.00
Integ.-Grad. [7]	67.4	72.8	73.2	1.15	82.5	82.5	85.3	<b>1.11</b>	59.4	58.3	58.3	0.98
SmoothGrad [8]	68.7	75.3	78.0	<b>1.20</b>	83.0	85.7	86.3	<b>1.13</b>	50.3	55.0	61.4	0.93
GradCAM [10]	77.6	85.7	84.1	<b>1.34</b>	81.9	83.5	82.4	1.10	54.4	52.5	54.1	0.90
Occlusion [18]	71.0	75.7	78.1	<b>1.22</b>	78.8	86.1	82.9	1.10	51.0	60.2	55.1	0.92
Grad.-Input [69]	65.8	63.3	67.9	1.06	76.5	82.9	79.5	1.05	50.0	57.6	62.6	0.95

Table 1: **Utility-K and Utility scores.** *Utility-K* scores across the 3 sessions for each attribution method, for each of the 3 datasets considered, followed by the *Utility* scores. Higher is better. The *Utility* scores of attribution methods that are statistically significant are **bolded**.

providing explanations along with training images simply increases participants’ engagement in the task. As we will show in Sec. 5, such non-informative explanations actually led to a decrease in participants’ ability to predict the classifier’s decisions – suggesting that giving a wrong explanation is worse than giving no explanations at all.

Each participant was only tested on a single condition to avoid possible experimental confounds. The main experiment was divided into 3 training sessions (with 5 training samples in each) each followed by a brief test. In each individual trial, an image was presented with the associated prediction of the model, either alone for the baseline condition or together with an explanation for the experimental and control condition. After a brief training phase (5 samples), participants’ ability to predict the classifier’s output was evaluated on 7 new samples during a test phase. During the test phase, no explanation was provided to limit confounding effects: one possible effect is if the explanation leaks information about the class label.<sup>‡</sup> We also propose to use a reservoir that subjects can refer to during the testing phase to minimize memory load as a confounding factor which was reported in [61] (see SI for an illustration).

**Datasets and models** We performed three distinct experiments in total – using a variety of neural network architectures and 6 representative attributions methods. Each of these experiments aimed at testing the usefulness of the explanation in a different context.

Our first scenario focuses on the detection of biases in AI systems using the popular Wolf vs. Husky dataset from [6] where an evaluation measure was already proposed around the usefulness of explanations for humans to detect biases. This makes it a good control experiment to measure the effectiveness of the framework proposed in Sec. 3. For this first experiment, we used the same model as in the original paper: InceptionV1 [70], and a similar dataset of Husky and Wolf images to bias the model. In this situation where prior knowledge of subjects can affect their Meta-predictor score, we balance data correctness (50% of correct/incorrect examples shown). Therefore, a subject relying only on their prior knowledge will end up as a bad Meta-predictor of the model. For this experiment, the results come from  $n = 242$  subjects who all passed our screening process.

In our second scenario, we focus on a representative challenging image categorization task which would be hard to solve by a non-expert untrained human participant and the goal is for the end-user to understand the strategy that was discovered by the AI system. Here, we chose the leaf dataset described in [48]. We selected 2 classes from this dataset (Betulaceae and Celastracea) that could not be classified by shape to reduce the chances that participants will discover the solution on their own – forcing them instead to rely on non-trivial features highlighted by the explanations (veins, leaf margin, etc). This scenario is far from being artificial as it reflects a genuine problem for the paleobotanist [49]. Can explainability methods help non-specialists discover the strategies discovered by an AI system? As participants are lay people from Amazon Mechanical Turk we do not expect them to be experts in botany, therefore we did not explicitly try to control for prior knowledge. In this experiment,  $n = 240$  subjects passed all our screening and other filtering processes.

<sup>‡</sup>Imagine an attribution method that would solely encode the classifiers’ prediction. Participants would be able to guess the classifier’s prediction perfectly from the explanation but the explanation per se would not help participants understand how the classifier works.

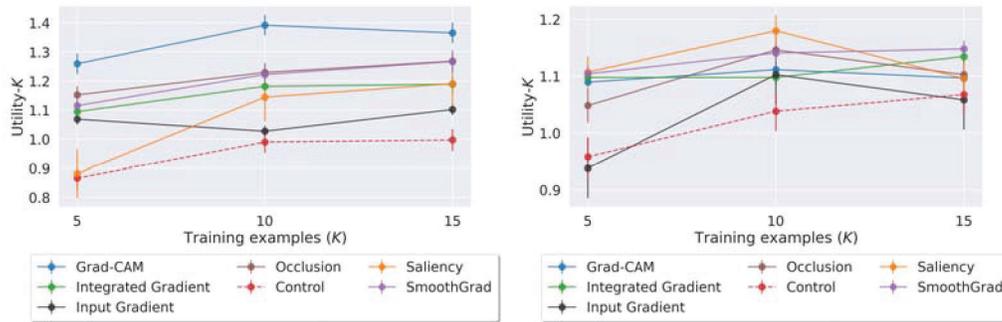


Figure 3: *Utility-K* for both Husky vs. Wolf (left) and the Leaves (right) dataset. The *Utility-K* of the explanation, or the accuracy of the human Meta-predictor after training, is measured after each training session (3 in total) for the scenario (1) of bias detection (on the left) and the scenario (2) concerning the identification of new strategies. Concerning the first scenario, all methods have a positive effect on the score obtained - they improve the subjects' ability to predict the model - and are thus useful to better understand the model. Grad-CAM, Occlusion and SmoothGrad are particularly useful for bias detection. On the Leaves dataset [48], explanations are also useful, but specifically Saliency, SmoothGrad and Integrated Gradients.

Finally, our last scenario focuses on identifying cases where an AI system fails<sup>§</sup> using ImageNet [50], also used in previous explainability work [12, 24, 30, 51, 59, 71]. We used this dataset because we expect it to be representative of real-world scenarios where it is difficult to understand what the model relies on for classification which makes it very difficult to understand these failure cases. Moreover, previous work has pointed out that attribution methods are not useful on this dataset [59], we have thus chosen to extend our analysis to this particular case. We use a ResNet50 [72] pretrained on this dataset as predictor. Because prior knowledge is a major confounding factor on ImageNet, we select a pair of classes that was heavily miss-classified by the model, to be able to show subjects 50% of correct/incorrect predictions: the pair Kit Fox and Red Fox fits this requirement. In this experiment, we analyzed data from  $n = 241$  participants who passed our screening and filtering processes.

For all experiments, we compared 6 representative attribution methods: Saliency (SA) [17], Gradient  $\odot$  Input (GI) [19], Integrated Gradients (IG) [7], Occlusion (OC) [18], SmoothGrad (SG) [8] and Grad-CAM (GC) [10]. Further information on these methods can be found in SI. Table 1 summarizes all the results from our psychophysics experiments.

## 5 Results

**Scenario (1): Bias detection** Fig. 3 shows the *Utility-K* scores for each method after different numbers of training samples were used to train participants for the biased dataset of Husky vs. Wolf. The *Utility* score encodes the quality of the explanations provided by a method, the higher the score, the better the method, with the baseline score being 1 (every score is divided by the baseline score corresponding to human accuracy after training without explanations).

A first observation is that the explanations have a positive effect on the *Utility-K* score: the explanation allows participants to better predict the model's decision (as the *Utility* scores are above 1). These results are consistent with those reported in [6]. This is confirmed with an Analysis of Variance (ANOVA) for which we found a significant main effect, with a medium effect size ( $F(7, 234) = 9.19$ ,  $p < .001$ ,  $\eta^2 = 0.089$ ). Moreover, the only score below the baseline is that of the control explanation, which do not make use of the model. We further explore our results by performing pairwise comparisons using Tukey's Honestly Significant Difference [73] to compare the different explanations against the baseline. We found 3 explainability methods to be significantly better than the baseline: Grad-CAM ( $p < 0.001$ ), Occlusion ( $p = 0.01$ ) and SmoothGrad ( $p = 0.034$ ). Thus, participants who received the Grad-CAM, Occlusion or SmoothGrad explanations performed much better than those who did not receive them.

<sup>§</sup>We acknowledge the existence of some overlap between the scenario 1 and scenario 3 as bias detection is a special case of a failure case. The reason we still use scenario 1 is because of the work previously done on it, allowing us to validate our framework.

**Scenario (2): Identifying an expert strategy** In Fig. 3, we show results on the Leaves dataset. An ANOVA analysis across all conditions revealed a significant main effect, albeit small ( $F(7, 232) = 4.29, p < .001, \eta^2 = 0.042$ ). This implies that explanation also had a positive effect resulting in better Meta-predictor in this use case. A Tukey’s Honestly Significant Difference test suggests that the best explanations are Saliency, SmoothGrad and Integrated Gradients as they are the only ones to be significantly better than our baseline (WE) ( $p = .004, p = .007$  and  $p = .03$  respectively). An interesting result is that SmoothGrad seems to be consistently useful across both use cases where explanations are indeed practically useful. A more surprising result is that Saliency which was one of the worst explanations for bias detection, is now the best explanation on this use case (We discuss possible reasons in SI).

**Scenario (3): Understanding failure cases** Table 1 shows that, on the ImageNet dataset, none of the methods tested exceeded baseline accuracy. Indeed, the experiment carried out, even with an improved experimental design, led us to the same conclusion as previous works [59]: none of the tested attribution methods are useful (ANOVA:  $F(7, 233) = 1.26, p > .05$ ). In the use case of understanding failure cases on ImageNet, **no attributions methods** seem to be useful.

### Why do attribution methods fail?

After studying the usefulness of attribution methods across 3 real-world scenarios for eXplainable AI, we found that attribution methods help, sometimes, but not always. We are interested in better understanding why sometimes attribution methods fail to help. Because this question has yet to be properly studied, there is no consensus if we can still make attribution methods work on those cases with incremental quantitative improvements. In the follow-up sections we explore 3 hypothesis to answer that question.

**Faithfulness as a proxy for Utility?** Faithfulness is often described as one of the key desiderata for a good explanation [23, 74, 75]. If an explanation fails to be sufficiently faithful, the rules it highlights won’t allow a user to understand the inner-working of the model. Thus, a lack of faithfulness on ImageNet could explain our results. To test this hypothesis, we use the *faithfulness* metrics: Deletion[9, 28], commonly used to compare attribution methods [9, 12, 24, 25, 28, 29]. A low Deletion score indicates a good *faithfulness*, thus for ease of reading we report the *faithfulness* score as  $1 - \text{Deletion}$  such that a higher *faithfulness* score is better.

Fig 4 shows the linear relationship between our Utility metrics and the faithfulness scores computed for every attribution method across all 3 datasets. We observe two main trends: 1) **There does not appear to be any specific pattern regarding faithfulness that could explain why attribution methods are not useful for ImageNet**, and 2) the least useful attribution methods for both use cases for which methods help (Bias and Leaves) are some of the leading methods in the field measured by the faithfulness metric. We also found a weak, if maybe anti-correlated, relation between faithfulness and usefulness: just focusing on making attribution methods more faithful does not translate to having methods with higher practical usefulness for end-users. And, in fact, focusing too heavily on faithfulness seem to come at the expense of usefulness, resulting in explanations that are counter-intuitively less useful. This second observation may seem rather alarming for the field given that the faithfulness measure is one of the driving benchmarks.

**Are explanations too complex?** Using the trade-off between completeness and understandability previously discussed in Section 3, we formulate another hypothesis: some explanations may be faithful but too complex and therefore cannot be understood by humans. In that view, an explanation with low complexity would tend to be more useful.

As a simple measure of the complexity of visual explanations, it would be ideal to be able to compute the Kolmogorov complexity [76] of each explanation. It was shown in previ-

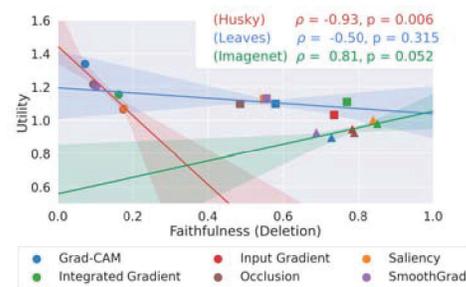


Figure 4: **Utility vs Faithfulness correlation.** The results suggest that current faithfulness metrics are poor predictors of the end-goal usefulness of explanation methods. Concerning the ImageNet dataset (triangle marker), the *Utility* scores are insignificant since none of the methods improves the baseline.

ous work [77] to correlate well with human-derived ratings for the complexity of natural images [78, 79]. As suggested by [76, 80] we used a standard compression technique (JPEG) to approximate the Kolmogorov complexity. Fig. 5 shows the *Utility* vs complexity score of attribution methods for each dataset. For one of the datasets where attribution methods help, the results suggest the presence of a strong correlation between *usefulness* and *complexity*: the least complex method is the most useful to end-users. For the other datasets, the results are either not conclusive (Leaves), or are not relevant as methods are not useful (ImageNet). Overall, **across datasets there is no significant difference in the complexity of explanations** that can explain why attribution methods do not help on ImageNet. This could be because the Kolmogorov Complexity does not perfectly reflect human visual complexity, or because this is not the key element to explain failure cases of attribution methods.

### An intrinsic limitation of Attribution methods?

The role of attribution methods is to help identify “where” to look in an image to understand the basis for a system’s prediction. However, attribution methods do not tell us “what” in those image regions is driving decisions. For categorization problems which involve perceptually similar classes (such as when discriminating between different breeds of dogs) and fine-grained categorization problems more generally, simply looking at diagnostic image regions tells the user very little about the specific shape property being relevant. For instance, knowing that the ear shape is being used for recognition does not say what specific shape feature is being encoded (e.g., pointed vs. round or narrow vs. broad base, etc). Our main hypothesis is that such a lack of explicit “what” information is precisely what is driving the failure of attribution methods on our ImageNet use-case.

To test this hypothesis, we estimated the perceptual similarity between classes measured within diagnostic regions (see SI for more details) using the Learned Perceptual Image Patch Similarity (LPIPS) metric [81] as it has been shown to approximate human perceptual similarity judgments well [81, 82]. We report the perceptual similarity score as  $1 - \text{LPIPS}$  score so that a high score means a high similarity. Fig. 6 shows the correlation between the perceptual similarity scores vs. our *Utility* scores on all methods and datasets studied. Our results suggest a strong correlation between perceptual similarity and practical usefulness: the more perceptually similar discriminative features of both classes are, the less useful attribution methods become. More importantly, the results across datasets show that on ImageNet, where attribution methods do not help, every method has a high similarity score. This result suggests that **after a certain threshold of perceptual similarity, attribution methods might no longer be useful, no matter how faithful or low in complexity the explanation is**. Overall, the results suggest that the perceptual similarity of discriminative features could explain why attribution methods fail on ImageNet.

## 6 Discussion

In summary, we conducted a large-scale human psychophysics experiment to test the utility of explainability methods in real-world scenarios. Our work shows that in two of the three tested scenarios (bias detection and identification of new strategies), explainability methods have indeed

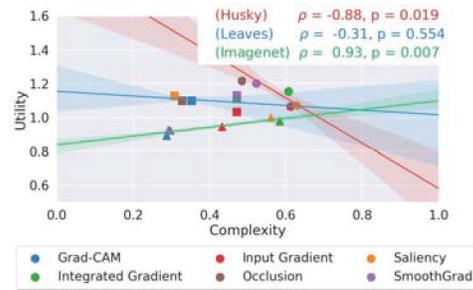


Figure 5: *Utility vs Complexity correlation*. This suggests a weak but possibly existing link between *Complexity* and *Utility* scores.

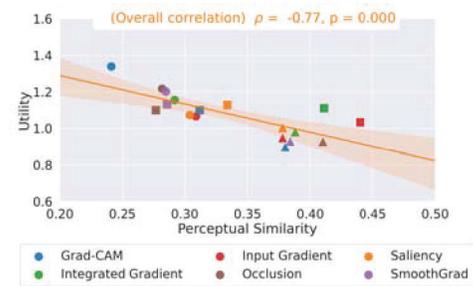


Figure 6: *Perceptual Similarity scores vs. Utility*. The perceptual similarity of highlighted regions by a given attribution method for both classes is measured, for each method, for each dataset. We observe a strong correlation between the perceptual similarity of features highlighted for both classes and the practical usefulness of methods.

progressed and they provide meaningful assistance to human end-users. Nevertheless, we identified a scenario (understanding failure case) for which none of the tested attribution methods were helpful. This result is consistent with previous work [59] and highlights a fundamental challenge for XAI.

Further analysis of associated faithfulness performance metrics driving the development of explainability methods revealed that they did not correlate with our empirical measure of utility – suggesting that they might not be suited anymore to move the field forward. We also investigated the possibility that the complexity of individual explanations may play a role in explaining human failures to learn to leverage those explanations to understand the model and, while we found a weak correlation between complexity and our empirical measure of utility, this correlation appears too low to explain the failure of these methods.

Finally, because attribution methods appear to be just as faithful and low in complexity whether they are useful or not, we explored the possibility that their failure lies, not in the quality of their explanations, but in the intrinsic limitations of attribution methods. If fully grasping the strategy of a model requires understanding, not just “where” to look (as revealed by attribution maps) but also “what” to look at, something not currently revealed by these methods, attribution methods will not help. Our assumption is that the need for finer “what” information should arise when diagnostic image locations across classes look perceptually very similar and potentially semantically related for certain classification problems (e.g., looking at the ears or the snout to discriminate between breeds of cats and dogs) and one needs to identify what visual features are driving decisions. We computed a perceptual score for classification problems by estimating the perceptual similarity between diagnostic image regions (as predicted by attribution methods) and found that, indeed, when this score predicts a certain level of perceptual similarity between classes, attribution methods fail to contribute useful information to human users, regardless of the faithfulness or complexity of the explanations. This suggests that explainability methods may need to communicate additional information to the end user beyond attribution maps.

## 7 Limitation and broader impact

**Limitations.** Our definition of the usefulness of an explanation is quite general. However, there are several limitations associated with our approach to estimate usefulness. First, our approach does not take into account the level of machine learning expertise of the user (the most useful explanation for a novice might not be the best one for an expert). Second, the need for human participants to evaluate explanations brings challenges compared with the automated metrics currently used in the field. However, by releasing all our data and software<sup>¶</sup>, we hope to encourage the adoption of the approach to evaluate future explainability methods and to assess the overall progress towards the development of more human-interpretable AI systems.

**Broader impacts.** While the increasing use of AI in real-world scenarios has shown potential to do good [46–48], it has also shown its potential for harm, especially when models rely on shortcuts (e.g., relying on spurious correlations in the training set that leads to unintended racial bias [39–41]). To identify such shortcuts, the field of Explainable AI has developed a lot of explainability methods; but it is not always clear which one performs best when trying to understand a model. We hope our evaluation method can help in this regard, assisting the AI practitioner in better identifying bias and shortcuts that may unfairly discriminate against groups of people.

## Acknowledgments

This work was conducted as part the DEEL<sup>¶</sup> project. It was funded by ONR grant (N00014-19-1-2029), NSF grant (IIS-1912280 and EAR-1925481), and the Artificial and Natural Intelligence Toulouse Institute (ANITI) grant #ANR19-PI3A-0004. The computing hardware was supported in part by NIH Office of the Director grant #S10OD025181 via the Center for Computation and Visualization (CCV) at Brown University. J.C. has been partially supported by funding from the Valencian Government (Conselleria d’Innovació, Universitats, Ciència i Societat Digital) by virtue of a 2022 grant agreement (convenio singular 2022).

<sup>¶</sup>[github.com/serre-lab/Meta-predictor](https://github.com/serre-lab/Meta-predictor)

<sup>¶</sup><https://www.deel.ai/>

## References

- [1] Franck Mamalet et al. *White Paper Machine Learning in Certified Systems*. 2021.
- [2] Robert Geirhos et al. “Shortcut Learning in Deep Neural Networks”. In: *arXiv preprint arXiv:2004.07780* (2020).
- [3] Alexander D’Amour et al. “Underspecification presents challenges for credibility in modern machine learning”. In: *arXiv preprint arXiv:2011.03395* (2020).
- [4] Sahar Shahamatdar et al. “Deceptive learning in histopathology”. In: *bioRxiv* (2022).
- [5] Thomas Fel et al. “Harmonizing the object recognition strategies of deep neural networks with humans”. In: *Advances in Neural Information Processing Systems (NeurIPS)* (2022).
- [6] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. ““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*. 2016.
- [7] Mukund Sundararajan, Ankur Taly, and Qiqi Yan. “Axiomatic Attribution for Deep Networks”. In: *Proceedings of the International Conference on Machine Learning (ICML)*. 2017.
- [8] Daniel Smilkov et al. “SmoothGrad: removing noise by adding noise”. In: *Workshop on Visualization for Deep Learning, Proceedings of the International Conference on Machine Learning (ICML)*. 2017.
- [9] Vitali Petsiuk, Abir Das, and Kate Saenko. “Rise: Randomized input sampling for explanation of black-box models”. In: *Proceedings of the British Machine Vision Conference (BMVC)*. 2018.
- [10] Ramprasaath R. Selvaraju et al. “Grad-CAM: Visual Explanations From Deep Networks via Gradient-Based Localization”. In: *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*. 2017.
- [11] Drew Linsley et al. “Learning what and where to attend”. In: *Proceedings of the International Conference on Learning Representations (ICLR)* (2018).
- [12] Fel Thomas et al. “Look at the Variance! Efficient Black-box Explanations with Sobol-based Sensitivity Analysis”. In: *Advances in Neural Information Processing Systems (NeurIPS)*. 2021.
- [13] Thomas Fel et al. “Don’t Lie to Me! Robust and Efficient Explainability with Verified Perturbation Analysis”. In: *ICML 2022, Workshop on Formal Verification of Machine Learning* (2022).
- [14] Paul Novello, Thomas Fel, and David Vigouroux. “Making Sense of Dependence: Efficient Black-box Explanations Using Dependence Measure”. In: *Advances in Neural Information Processing Systems (NeurIPS)* (2022).
- [15] Éloi Zablocki et al. “Explainability of vision-based autonomous driving systems: Review and challenges”. In: *arXiv preprint arXiv:2101.05307* (2021).
- [16] Bryce Goodman and Seth Flaxman. “European Union regulations on algorithmic decision-making and a “right to explanation””. In: *AI magazine* 38.3 (2017), pp. 50–57.
- [17] Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. “Deep inside convolutional networks: Visualising image classification models and saliency maps”. In: *Workshop Proceedings of the International Conference on Learning Representations (ICLR)*. 2014.
- [18] Matthew D Zeiler and Rob Fergus. “Visualizing and understanding convolutional networks”. In: *Proceedings of the IEEE European Conference on Computer Vision (ECCV)*. 2014.
- [19] Marco Ancona et al. “Towards better understanding of gradient-based attribution methods for Deep Neural Networks”. In: *Proceedings of the International Conference on Learning Representations (ICLR)*. 2018.
- [20] Zachary C. Lipton. “The Mythos of Model Interpretability”. In: *Workshop on Human Interpretability in Machine Learning, ICML*. 2016.
- [21] Diogo V Carvalho, Eduardo M Pereira, and Jaime S Cardoso. “Machine learning interpretability: A survey on methods and metrics”. In: *Electronics* 8.8 (2019), p. 832.
- [22] Alejandro Barredo Arrieta et al. “Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI”. In: *Information Fusion* 58 (2020), pp. 82–115.
- [23] Thomas Fel and David Vigouroux. “Representativity and Consistency Measures for Deep Neural Network Explanations”. In: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*. 2022.
- [24] Ruth C. Fong and Andrea Vedaldi. “Interpretable Explanations of Black Boxes by Meaningful Perturbation”. In: *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*. 2017.
- [25] Ruth Fong, Mandela Patrick, and Andrea Vedaldi. “Understanding deep networks via extremal perturbations and smooth masks”. In: *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*. 2019.
- [26] Andrew Elliott, Stephen Law, and Chris Russell. “Explaining Classifiers using Adversarial Perturbations on the Perceptual Ball”. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2021.

- [27] Jianming Zhang et al. “Top-down neural attention by excitation backprop”. In: *International Journal of Computer Vision* 126.10 (2018), pp. 1084–1102.
- [28] Wojciech Samek et al. “Evaluating the visualization of what a Deep Neural Network has learned”. In: *IEEE Transactions on Neural Networks and Learning Systems (TNNLS)*. 2015.
- [29] Andrei Kapishnikov et al. “Xrai: Better attributions through regions”. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2019.
- [30] Giang Nguyen, Daeyoung Kim, and Anh Nguyen. “The effectiveness of feature attribution methods and its correlation with automatic evaluation scores”. In: *Advances in Neural Information Processing Systems (NeurIPS)* (2021).
- [31] Finale Doshi-Velez and Been Kim. “Towards a rigorous science of interpretable machine learning”. In: *ArXiv e-print* (2017).
- [32] Tania Lombrozo and Susan Carey. “Functional explanation and the function of explanation”. In: *Cognition* 99.2 (2006), pp. 167–204.
- [33] Tania Lombrozo and Nicholas Z Gwynne. “Explanation and inference: Mechanistic and functional explanations guide property generalization”. In: *Frontiers in Human Neuroscience* 8 (2014), p. 700.
- [34] Ny Vasil, Azzurra Ruggeri, and Tania Lombrozo. “When and how children use explanations to guide generalizations”. In: *Cognitive Development* 61 (2022), p. 101144.
- [35] Joseph J Williams and Tania Lombrozo. “The role of explanation in discovery and generalization: Evidence from category learning”. In: *Cognitive science* 34.5 (2010), pp. 776–806.
- [36] Tania Lombrozo. “The instrumental value of explanations”. In: *Philosophy Compass* 6.8 (2011), pp. 539–551.
- [37] Tania Lombrozo. “Explanatory preferences shape learning and inference”. In: *Trends in Cognitive Sciences* 20.10 (2016), pp. 748–759.
- [38] Emily G Liquin and Tania Lombrozo. “Motivated to learn: An account of explanatory satisfaction”. In: *Cognitive Psychology* 132 (2022), p. 101453.
- [39] Julia Angwin et al. *Machine Bias*. 2016. URL: <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>.
- [40] Ziad Obermeyer et al. “Dissecting racial bias in an algorithm used to manage the health of populations”. In: *Science* 366.6464 (2019), pp. 447–453.
- [41] Abdallah Hussein Sham et al. “Ethical AI in facial expression analysis: racial bias”. In: *Signal, Image and Video Processing* (2022), pp. 1–8.
- [42] Thomas McGrath et al. “Acquisition of Chess Knowledge in AlphaZero”. In: *arXiv preprint arXiv:2111.09259* (2021).
- [43] Jonas Andriulis et al. “Domain-Level Explainability—A Challenge for Creating Trust in Superhuman AI Strategies”. In: *arXiv preprint arXiv:2011.06665* (2020).
- [44] David Silver et al. “Mastering the game of go without human knowledge”. In: *nature* 550.7676 (2017), pp. 354–359.
- [45] Jonas Degraeve et al. “Magnetic control of tokamak plasmas through deep reinforcement learning”. In: *Nature* 602.7897 (2022), pp. 414–419.
- [46] John Jumper et al. “Highly accurate protein structure prediction with AlphaFold”. In: *Nature* 596.7873 (2021), pp. 583–589.
- [47] Alex Davies et al. “Advancing mathematics by guiding human intuition with AI”. In: *Nature* 600.7887 (2021), pp. 70–74.
- [48] Peter Wilf et al. “Computer vision cracks the leaf code”. In: *Proceedings of the National Academy of Sciences* 113.12 (2016), pp. 3305–3310.
- [49] Edward J Spagnuolo, Peter Wilf, and Thomas Serre. “Decoding family-level features for modern and fossil leaves from computer-vision heat maps”. In: *American journal of botany* (2022).
- [50] J. Deng et al. “ImageNet: A Large-Scale Hierarchical Image Database”. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2009.
- [51] Sara Hooker et al. “A Benchmark for Interpretability Methods in Deep Neural Networks”. In: *Advances in Neural Information Processing Systems (NIPS)*. 2019.
- [52] Leila Arras et al. “Explaining recurrent neural network predictions in sentiment analysis”. In: *Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis (WASSA) in ENLP*. 2017.
- [53] Leila Arras et al. ““What is relevant in a text document?”: An interpretable machine learning approach”. In: *PloS one* 12.8 (2017), e0181142.
- [54] Cheng-Yu Hsieh et al. “Evaluations and methods for explanation through robustness analysis”. In: *Proceedings of the International Conference on Learning Representations (ICLR)* (2020).
- [55] Oisín Mac Aodha et al. “Teaching categories to human learners with visual explanations”. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2018.

- [56] Arjun Chandrasekaran et al. “Do explanations make VQA models more predictable to a human?” In: *arXiv preprint arXiv:1810.12366* (2018).
- [57] Yasmeeen Alufaisan et al. “Does Explainable Artificial Intelligence Improve Human Decision-Making?” In: *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*. 2021.
- [58] Felix Biessmann and Dionysius Refiano. “Quality Metrics for Transparent Machine Learning With and Without Humans In the Loop Are Not Correlated”. In: *Proceedings of the ICML Workshop on Theoretical Foundations, Criticism, and Application Trends of Explainable AI held in conjunction with the 38th International Conference on Machine Learning (ICML)*, 2021.
- [59] Hua Shen and Ting-Hao Huang. “How Useful Are the Machine-Generated Interpretations to General Users? A Human Evaluation on Guessing the Incorrectly Predicted Labels”. In: *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*. Vol. 8. 1. 2020, pp. 168–172.
- [60] Dong Nguyen. “Comparing automatic and human evaluation of local explanations for text classification”. In: *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*. 2018, pp. 1069–1078.
- [61] Peter Hase and Mohit Bansal. “Evaluating Explainable AI: Which Algorithmic Explanations Help Users Predict Model Behavior?” In: *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL Short Papers)*. 2020.
- [62] Giang Nguyen, Mohammad Reza Taesiri, and Anh Nguyen. “Visual correspondence-based explanations improve AI robustness and human-AI team accuracy”. In: *Advances in Neural Information Processing Systems (NeurIPS)* (2022).
- [63] Sunnie S. Y. Kim et al. “HIVE: Evaluating the Human Interpretability of Visual Explanations”. In: *European Conference on Computer Vision (ECCV)*. 2022.
- [64] Leon Sixt et al. “Do Users Benefit From Interpretable Vision? A User Study, Baseline, And Dataset”. In: *Proceedings of the International Conference on Learning Representations (ICLR)* (2022).
- [65] Been Kim, Rajiv Khanna, and Oluwasanmi O Koyejo. “Examples are not enough, learn to criticize! criticism for interpretability”. In: *Advances in Neural Information Processing Systems (NeurIPS)* (2016).
- [66] Danish Pruthi et al. “Evaluating Explanations: How much do explanations from the teacher aid students?” In: *ArXiv e-print*. 2020.
- [67] Alon Jacovi and Yoav Goldberg. “Towards faithfully interpretable NLP systems: How should we define and evaluate faithfulness?” In: *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL Short Papers)* (2020).
- [68] L. Itti. “Models of bottom-up attention and saliency”. In: *Neurobiology of attention* (2005).
- [69] Peyton Greenside Avanti Shrikumar and Anshul Kundaje. “Learning Important Features Through Propagating Activation Differences”. In: *Proceedings of the 34th International Conference on Machine Learning (ICML)*. 2017.
- [70] Christian Szegedy et al. “Going deeper with convolutions”. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2015.
- [71] Andrew Elliott, Stephen Law, and Chris Russell. “Explaining Classifiers using Adversarial Perturbations on the Perceptual Ball”. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2021, pp. 10693–10702.
- [72] Kaiming He et al. “Deep residual learning for image recognition”. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2016.
- [73] John W Tukey. “Comparing individual means in the analysis of variance”. In: *Biometrics* (1949).
- [74] Umang Bhatt, Adrian Weller, and José M. F. Moura. “Evaluating and Aggregating Feature-based Model Explanations”. In: *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*. 2020.
- [75] Chih-Kuan Yeh et al. “On the (In)Fidelity and Sensitivity for Explanations”. In: *Advances in Neural Information Processing Systems (NIPS)*. 2019.
- [76] Ming Li et al. “The similarity metric”. In: *IEEE transactions on Information Theory* 50.12 (2004), pp. 3250–3264.
- [77] Matthieur Perreira Da Silva, Vincent Courboulay, and Pascal Estraillier. “Image complexity measure based on visual attention”. In: *2011 18th IEEE International Conference on Image Processing*. IEEE. 2011, pp. 3281–3284.
- [78] Alex Forsythe, Gerry Mulhern, and Martin Sawey. “Confounds in pictorial sets: The role of complexity and familiarity in basic-level picture processing”. In: *Behavior research methods* 40.1 (2008), pp. 116–129.
- [79] Alexandra Forsythe. “Visual complexity: is that all there is?” In: *International Conference on Engineering Psychology and Cognitive Ergonomics*. Springer. 2009, pp. 158–166.
- [80] Steven de Rooij and Paul Vitányi. “Approximating rate-distortion graphs of individual data: Experiments in lossy compression and denoising”. In: *arXiv preprint cs/0609121* (2006).

- [81] Richard Zhang et al. “The unreasonable effectiveness of deep features as a perceptual metric”. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2018, pp. 586–595.
- [82] Vedant Nanda et al. “Exploring Alignment of Representations with Human Perception”. In: *arXiv preprint arXiv:2111.14726* (2021).