

Explaining Deep Neural Networks and Beyond: A Review of Methods and Applications

Wojciech Samek^{†*}, *Member, IEEE*, Grégoire Montavon^{†*}, Sebastian Lapuschkin, Christopher J. Anders,
and Klaus-Robert Müller*, *Member, IEEE*

Abstract—With the broader and highly successful usage of machine learning in industry and the sciences, there has been a growing demand for Explainable AI. Interpretability and explanation methods for gaining a better understanding about the problem solving abilities and strategies of nonlinear Machine Learning, in particular, deep neural networks, are therefore receiving increased attention. In this work we aim to (1) provide a timely overview of this active emerging field, with a focus on ‘post-hoc’ explanations, and explain its theoretical foundations, (2) put interpretability algorithms to a test both from a theory and comparative evaluation perspective using extensive simulations, (3) outline best practice aspects i.e. how to best include interpretation methods into the standard usage of machine learning and (4) demonstrate successful usage of explainable AI in a representative selection of application scenarios. Finally, we discuss challenges and possible future directions of this exciting foundational field of machine learning.

Index Terms—Interpretability, deep learning, neural networks, black-box models, explainable artificial intelligence (XAI), model transparency.

I. INTRODUCTION

A main goal of machine learning is to learn accurate decision systems respectively predictors that can help automatizing tasks, that would otherwise have to be done by humans. Machine Learning (ML) has supplied a wealth of algorithms that have demonstrated important successes in the sciences and industry; most popular ML work horses are considered to be kernel methods (e.g. [190], [164], [132], [163], [194]) and particularly during the last decade deep learning methods (e.g. [23], [52], [108], [107], [161], [70]) have gained highest popularity.

As ML is increasingly used in real-world applications, a general consensus has emerged that high prediction accuracy

alone may not be sufficient in practice [104], [28], [159]. Instead, in practical engineering of systems, critical features that are typically considered beyond excellent prediction itself are (a) robustness of the system to measurement artefacts or adversarial perturbations [182], (b) resilience to drifting data distributions [49], (c) ability to accurately assess the confidence of its own predictions [139], [135], (d) safety and security aspects [21], [84], [26], [193], (e) legal requirements or adherence to social norms [54], [60], (f) ability to complement human expertise in decision making [82], or (g) ability to reveal to the user the interesting correlations it has found in the data [88], [165].

Orthogonal to the quest for better and more holistic machine learning models, Explainable AI (XAI) [159], [73], [112], [128], [18] has developed as a subfield of machine learning that seeks to augment the training process, the learned representations and the decisions with human-interpretable explanations. An example is *medical diagnosis*, where the input examples (e.g. histopathological images) come with various artifacts (e.g. stemming from image quality or suboptimal annotations) that have in principle nothing to do with the diagnostic task, yet, due to the limited amount of available data, the ML model may harvest specifically these spurious correlations with the prediction target (e.g. [61], [177]). Here interpretability could point at anomalous or awkward decision strategies before harm is caused in a later usage as a diagnostic tool.

Similarly essential when using ML in the sciences is again interpretability, since ideally, the transparent ML model — having learned from data — may have embodied scientific knowledge that would subsequently provide insight to the scientist, occasionally this can even be novel scientific insight (see e.g. [165]). — Note that in numerous scientific applications it has been most common so far to use linear models [151], favoring interpretability often at the expense of predictivity (see e.g. [63], [117]).

To summarize, there is a strong push toward better understanding ML systems that are being used and in consequence blackbox algorithms are more and more abandoned for many applications. This growing consensus has led to a strong growth of a subfield of ML, namely explainable AI (XAI) that strives to produce transparent nonlinear learning methods, and supplies novel theoretical perspectives on machine learning models, along with powerful practical tools for a better understanding and interpretation of AI systems.

In this review paper, we will summarize the recent exciting developments, present different classes of XAI methods that

[†] W. Samek and G. Montavon contributed equally to this work.

* Corresponding authors: W. Samek, G. Montavon and K.-R. Müller.

W. Samek is with the Dept. of Artificial Intelligence, Fraunhofer Heinrich Hertz Institute, 10587 Berlin, Germany, and with BIFOLD – Berlin Institute for the Foundations of Learning and Data, 10587 Berlin, Germany. (e-mail: wojciech.samek@hhi.fraunhofer.de).

G. Montavon and C. Anders are with the Machine Learning Group, Technische Universität Berlin, 10587 Berlin, Germany, and with BIFOLD – Berlin Institute for the Foundations of Learning and Data, 10587 Berlin, Germany. (e-mail: gregoire.montavon@tu-berlin.de).

S. Lapuschkin is with the Dept. of Artificial Intelligence, Fraunhofer Heinrich Hertz Institute, 10587 Berlin, Germany.

K.-R. Müller is with the Machine Learning Group, Technische Universität Berlin, 10587 Berlin, Germany, and also with the Dept. of Artificial Intelligence, Korea University, Seoul 136-713, South Korea, the Max Planck Institute for Informatics, 66123 Saarbrücken, Germany, and BIFOLD – Berlin Institute for the Foundations of Learning and Data, 10587 Berlin, Germany. (e-mail: klaus-robert.mueller@tu-berlin.de).

Manuscript is currently under review at PROCEEDINGS OF THE IEEE.

have been proposed in the context of deep neural networks, provide theoretical insights, and highlight the current best practices when applying these methods. Note finally, that we do not attempt an encyclopedic treatment of all available XAI literature, rather, we present a slightly biased point of view (and in doing so we often draw from the work of the authors). In particular, we focus on *post-hoc explanation methods*, which take any model, typically the best performing one, and analyze it in a second step in order to uncover its decision strategy. We also provide — to the best of our knowledge — reference to other related works for further reading.

II. TOWARDS EXPLAINING DEEP NEURAL NETWORKS

Before discussing aspects of the problem of explanation that are specific to deep neural networks, we first introduce some basics of Explainable AI, which apply to a fairly general class of machine learning models. The ML model will be assumed to be *already trained* and the input-output relation it implements will be abstracted by some *function*:

$$f: \mathbb{R}^d \rightarrow \mathbb{R}.$$

This function receives as input a vector of real-valued features $\mathbf{x} = (x_1, \dots, x_d)$ typically corresponding to various sensor measurements. The function produces as an output a real-valued score on which the decision is based. A sketch of such function receiving two features x_1 and x_2 as input is given in Fig. 1.




Fig. 1. Example of a nonlinear function of the input features, which produces some prediction. The function can be approximated locally as a linear model.

In the context of ML classification, the function output can be interpreted as the amount of evidence for/against deciding in favor of a certain class. A classification decision is then obtained from the output score by testing whether the latter is above a certain threshold, or for multiclass problems, larger than the output score of other functions representing the remaining classes.

In a medical scenario, the function may receive as input an array of clinical variables, and the output of the function may represent the evidence for a certain medical condition [96]. In an engineering setting, the input could be the composition of some compound material, and the output could be a prediction of its strength [197] or stability.

Suppose a given instance is predicted by the machine learning model to be healthy, or a compound material is predicted to have high strength. We may choose to trust the prediction and go ahead with next step within an application scenario. However, we may benefit from taking a closer look

at that prediction, e.g. to verify that the prediction ‘healthy’ is associated to relevant clinical information, and not some spurious features that accidentally correlate with the predicted quantity in the dataset [101], [104]. Such problem can often be identified by building an *explanation* of the ML prediction [104].

Conversely, suppose that another instance is predicted by the machine learning model to be of low health or low strength. Here, an explanation could prove equally useful as it could hint at actions to be taken on the sample to improve its predicted score [189], e.g. possible therapies in a medical setting, or small adjustments of the compound design that lead to higher strength.

A. How to Explain: Global vs. Local

Numerous approaches have emerged to shed light onto machine learning predictions. Certain approaches such as activation-maximization [174], [136], [134] aim at a *global* interpretation of the model, by identifying prototypical cases for the output quantity

$$\mathbf{x}^* = \arg \max_{\mathbf{x}} f(\mathbf{x})$$

and allowing in principle to verify that the function has a high value only for the valid cases. While these prototypical cases may be interesting per se, both for model validation or knowledge discovery, such prototypes will be of little use to understand for a given example \mathbf{x} (say, near the decision boundary) what features play in favor or against the model output $f(\mathbf{x})$.

Specifically, we would like to know for that very example what input features contribute positively or negatively to the given prediction. These *local* analyses of the decision function have received growing attention and many approaches have been proposed [14], [201], [13], [147], [183]. For simple models with limited nonlinearity, the decision function can be approximated locally as the linear function [13]:

$$f(\mathbf{x}) \approx \sum_{i=1}^d \underbrace{[\nabla f(\tilde{\mathbf{x}})]_i}_{R_i} (x_i - \tilde{x}_i) \quad (1)$$

where $\tilde{\mathbf{x}}$ is some nearby root point (cf. Fig. 1). This expansion takes the form of a weighted sum over the input features, where the summand R_i can be interpreted as the contribution of feature i to the prediction. Specifically, an inspection of the summands reveals that a feature x_i will be attributed strong relevance if the following two conditions are met: (1) the feature must be expressed in the data, i.e. it differs from the reference value \tilde{x}_i , and (2) the model output should be sensitive to the presence of that feature, i.e. $[\nabla f(\tilde{\mathbf{x}})]_i \neq 0$. An explanation for the prediction can then be formed by the vector of relevance scores $(R_i)_i$. It can be given to the user as a histogram over the input features or as a heatmap.

For illustration, consider the problem of explaining a prediction for a data point from the Concrete Compressive Strength dataset [197]. For this data point, a simple two-layer neural network model predicts a low compressive strength. Applying




Fig. 2. Input example predicted to have low compressive strength, and a feature-wise explanation of the prediction. Red and blue color indicate positive and negative contributions.

the analysis above gives an explanation for this prediction, which we show in Fig. 2.

For this example low cement concentration and below average age are factors of low compressive strength, although this is partly compensated by a high quantity of blast furnace slag.

Furthermore, for an explanation to be interpretable by its receiver, the latter must be able to make sense of the input features. Some features such as ‘cement’, ‘water’, and ‘age’, are understandable to everyone, however, more technical terms such as ‘blast furnace slag’ or ‘superplasticizer’ may only be accessible to a domain expert. Therefore, when using these explanation techniques, we make the implicit assumption that those input features are interpretable to the receiver.

B. Deep Networks and the Difficulty of Explaining Them

In practice, linear models or shallow neural networks may not be sufficiently expressive to predict the task optimally. Deep neural networks have been proposed as a way of producing more predictive models. They can be abstracted as a sequence of layers

$$f(x) = f_L \circ \dots \circ f_1(x),$$

where each layer applies a linear transformation followed by an element-wise nonlinearity. Combining a large number of these layers endows the model with high prediction power. DNNs have proven especially successful on computer vision tasks [99], [175], [65]. However, DNN models are also much more complex and nonlinear, and quantities entering into the simple explanation model of Eq. (1) become considerably harder to compute and to estimate reliably.

A first difficulty comes from the multiscale and distributed nature of neural network representations. Some neurons are activated for only a few data points, whereas others apply more globally. The prediction is thus a sum of local and global effects, which makes it difficult (or impossible) to find a root point \tilde{x} that linearly expands to the prediction for the data point of interest. The transition from the global to local effect indeed introduces a nonlinearity, which Eq. (1) cannot capture.

A second source of instability arises from the high depth of recent neural networks, where a ‘shattered gradient’ effect was observed [16], noting that the gradient locally resembles white noise. In particular, it can be shown that for deep rectifier networks, the number of discontinuities of the gradient

can grow in the worst case exponentially with depth [129]. The shattered gradient effect is illustrated in Fig. 3 (left) for the well-established VGG-16 network [175]: The network is fed multiple consecutive video frames of an athlete lifting a barbell, and we observe the prediction for the output neuron ‘barbell’. The gradient of the prediction is changing its value much more quickly than the prediction itself. An explanation based on such gradient would therefore inherit this noise.

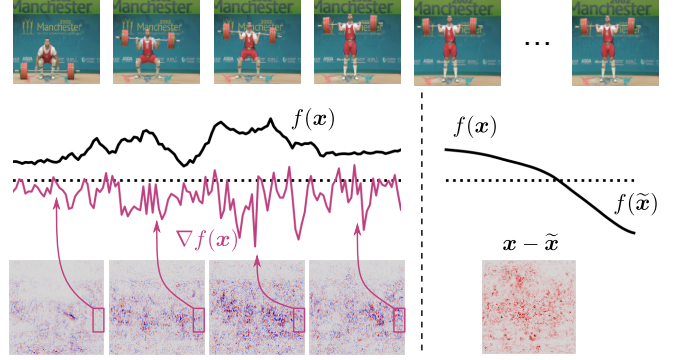


Fig. 3. Two difficulties encountered with DNNs. *Left*: Shattered gradient effect causing gradients to be highly varying and too noisy to be used for explanation. *Right*: Pathological minima in the function, making it difficult to search for meaningful reference points.

A last difficulty comes from the challenge of searching for a root point \tilde{x} on which to base the explanation, that is both close to the data and not an adversarial example [53], [135]. The problem is illustrated in Fig. 3 (right), where we showcase a reference point \tilde{x} that does not carry any meaningful visual difference to the original data x , but for which function output has changed dramatically. The problem of adversarial examples can be explained by the gradient noise, that causes the model to ‘overreact’ to certain pixel-wise perturbations, and also by the high dimensionality of the data ($224 \times 224 = 50176$ pixels for the ImageNet dataset) where many small pixel-wise effects cumulate into a large effect on the model output.

III. PRACTICAL METHODS FOR EXPLAINING DNNs

In view of the multiple challenges posed by analyzing deep neural network functions, building robust and practical methods to explain their decisions has developed into an own research area [128], [59], [159] and an abundance of methods have been proposed. In parallel, efficient software (cf. Appendix C for a list) makes these newly developed methods readily usable in practice, and allows researchers to perform systematic comparisons between them on reference models and datasets.

In this section, we focus on four families of explanation techniques: Interpretable Local Surrogates, Occlusion Analysis, Gradient-based techniques, and Layer-Wise Relevance Propagation. In our view, these techniques exemplify the current diversity of possible approaches to explaining predictions in terms of input features, and taken together provide a broad coverage of the types of models to explain and the practical use cases. We give references to further related methods in the

corresponding subsections. Table III in Appendix C provides a glossary of all referenced methods.

A. Interpretable Local Surrogates [147]

This category of methods aims to replace the decision function by a local surrogate model that is structured in a way that it is self-explanatory (an example of a self-explanatory model is the linear model). This approach is embodied in the LIME algorithm [147], which was successfully applied to DNN classifiers for images and text. Explanation can be achieved by first defining some local distribution $p_{\mathbf{x}}(\boldsymbol{\xi})$ around our data point \mathbf{x} , learning the parameter \mathbf{v} of the linear model that best matches the function locally:

$$\min_{\mathbf{v}} \int [f(\boldsymbol{\xi}) - \mathbf{v}^\top \boldsymbol{\xi}]^2 \cdot dp_{\mathbf{x}}(\boldsymbol{\xi})$$

and then extracting local feature contributions, e.g. $R_i = v_i x_i$. Because the method does not rely on the gradient of the original DNN model, it avoids some of the difficulties discussed in Section II-B. The LIME method also covers the incorporation of sparsity or simple decision tree to the surrogate model to further enhance interpretability. Additionally, the learned surrogate model may be based on its own set of *interpretable features*, allowing to produce explanations in terms of features that are maximally interpretable to the human. Other methods that explain by building a local surrogate include LORE [58] and Anchors [148]. Furthermore, a broader set of methods do not consider a specific location in the input space and build instead a global surrogate model of the decision function, where the surrogate model readily incorporates interpretability structures. We discuss these global ‘self-explainable’ models in Section III-E.

B. Occlusion Analysis [201]

Occlusion Analysis is a particular type of perturbation analysis where we repeatedly test the effect on the neural network output, of occluding patches or individual features in the input image [201], [208]:

$$R_i = f(\mathbf{x}) - f(\mathbf{x} \odot (1 - \mathbf{m}_i))$$

where \mathbf{m}_i is an indicator function for the patch or feature to remove, and ‘ \odot ’ denotes the element-wise product. A heatmap $(R_i)_i$ can be built from these scores highlighting locations where the occlusion has caused the strongest decrease of the function. Because occlusion may produce visual artefacts, inpainting occluded patterns (e.g. using a generative model [2]) rather than setting them to gray was proposed as an enhancement.

Attribution based on Shapley values [116], [179] (see Section V-A for a definition), can also be seen as an occlusion analysis. Here, instead of occluding features one at a time, a much broader set of occlusion patterns are considered, and this has the effect of also integrating global effects in the explanation. SHAP and Kernel SHAP [116] are practical algorithms to approximate Shapley values, that sample a few occlusion according to the probability distribution used to compute Shapley values, and then fit a linear surrogate model

that correctly predicts the effect of these occlusions on the output. An explanation can then be easily extracted, and this explanation retains some similarity with the original Shapley values.

A further extension of occlusion analysis is Meaningful Perturbation [47], where an occluding pattern is synthesized, subject to a sparsity constraint, in order to engender the maximum drop of the function value f . The explanation is then readily given by the synthesized pattern. The perturbation-based approach was latter embedded in a rate distortion theoretical framework [118].

C. Integrated Gradients / SmoothGrad [183], [176]

Integrated Gradients [183] explains by integrating the gradient $\nabla f(\mathbf{x})$ along some trajectory in input space connecting some root point $\tilde{\mathbf{x}}$ to the data point \mathbf{x} . The integration process addresses the problem of locality of the gradient information (cf. Section II-B), making it well-suited for explaining functions that have multiple scales. In the simplest form, the trajectory is chosen to be the segment $[\tilde{\mathbf{x}}, \mathbf{x}]$ connecting some root point to the data. Integrated gradients defines feature-wise scores as:

$$R_i(\mathbf{x}) = (x_i - \tilde{x}_i) \cdot \int_0^1 [\nabla f(\tilde{\mathbf{x}} + t \cdot (\mathbf{x} - \tilde{\mathbf{x}}))]_i dt$$

It can be shown that these scores satisfy $\sum_i R_i(\mathbf{x}) = f(\mathbf{x})$ and thus constitute a *complete* explanation. If necessary, the method can be easily extended to any trajectories in input space. For implementation purposes, integrated gradients must be discretized. Specifically, the continuous trajectory is approximated by a sequence of data points $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(N)}$. Integrated gradients is then implemented as shown in Algorithm 1.

Algorithm 1 Integrated Gradients

```

R = 0
for  $n = 1$  to  $N - 1$  do
   $\mathbf{R} = \mathbf{R} + \nabla f(\mathbf{x}^{(n)}) \odot (\mathbf{x}^{(n+1)} - \mathbf{x}^{(n)})$ 
end for
return  $\mathbf{R}$ 

```

The gradient can easily be computed using automatic differentiation. The operation ‘ \odot ’ denotes the element-wise product. The larger the number of discretization steps, the closer the output gets to the integral form, but the more computationally expensive the procedure gets.

Another popular gradient-based explanation method is *SmoothGrad* [176]. The function’s gradient is averaged over a large number of locations corresponding to small random perturbations of the original data point \mathbf{x} :

$$\nabla_{\text{smooth}} f(\mathbf{x}) = \mathbb{E}_{\boldsymbol{\varepsilon} \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})} [\nabla f(\mathbf{x} + \boldsymbol{\varepsilon})]$$

Like the method’s name suggests, the averaging process ‘smoothes’ the explanation, and in turn also addresses the shattered gradient problem described in Section II-B. (See also [130], [14], [174] for earlier gradient-based explanation techniques).

In Section IV, we experiment with a combination of Integrated Gradients and SmoothGrad [176], similar to Expected Gradients (cf. [181]), where relevance scores obtained from Integrated Gradients are averaged over several integration paths that are drawn from some random distribution. The resulting method preserves the advantages of Integrated Gradients and further reduces the gradient noise.

D. Layer-Wise Relevance Propagation [13]

The *Layer-wise Relevance Propagation* (LRP) method [13] makes explicit use of the layered structure of the neural network and operates in an iterative manner to produce the explanation. Consider the neural network

$$f(\mathbf{x}) = f_L \circ \dots \circ f_1(\mathbf{x})$$

First, activations at each layer of the neural network are computed until we reach the output layer. The activation score in the output layer forms the prediction. Then, a reverse propagation pass is applied, where the output score is progressively redistributed, layer after layer, until the input variables are reached. The redistribution process follows a conservation principle analogous to Kirchoff's laws in electrical circuits. Specifically, all 'relevance' that flows into a neuron at a given layer flows out towards the neurons of the layer below. At a high level, the LRP procedure can be implemented as a forward-backward loop, as shown in Algorithm 2.

Algorithm 2 Layer-wise Relevance Propagation

```

 $\mathbf{a}^{(0)} = \mathbf{x}$ 
for  $l = 1 \dots L$  do
   $\mathbf{a}^{(l)} = f_l(\mathbf{a}^{(l-1)})$ 
end for
 $\mathbf{R}^{(L)} = \mathbf{a}^{(L)}$ 
for  $l = L \dots 1$  do
   $\mathbf{R}^{(l-1)} = \text{relprop}(\mathbf{a}^{(l-1)}, \mathbf{R}^{(l)}, f_l)$ 
end for
return  $\mathbf{R}^{(0)}$ 

```

The function `relprop` performs redistribution from one layer to the layer below and is based on 'propagation rules' defining the exact redistribution policy. Examples of propagation rules are given later in this section, and their implementation is provided in Appendix B. The LRP procedure is shown graphically in Fig. 4.

While LRP can in principle be performed in any forward computational graph, a class of neural networks which is often encountered in practice, and for which LRP comes with efficient propagation rules that can be theoretically justified (cf. Section V) is deep rectifier networks [51]. The latter can be in large part abstracted as an interconnection of neurons of the type:

$$a_k = \max(0, \sum_{0,j} a_j w_{jk}),$$

where a_j denotes some input activation, and w_{jk} is the weight connecting neuron j to neuron k in the layer above. The notation $\sum_{0,j}$ indicates that we sum over all neurons j in the lower layer plus a bias term w_{0k} with $a_0 = 1$. For this class of




Fig. 4. Illustration of the LRP propagation procedure applied to a neural network. The prediction at the output is propagated backward in the network, using various propagation rules, until the input features are reached. The propagation flow is shown in red.

networks, various propagation rules have been proposed (cf. Fig. 4). For example, the LRP- γ rule [126] defined as

$$R_j = \sum_k \frac{a_j(w_{jk} + \gamma w_{jk}^+)}{\sum_{0,j} a_j(w_{jk} + \gamma w_{jk}^+)} R_k \quad (2)$$

redistributes based on the contribution of lower-layer neurons to the given neuron activation, with a preference for positive contributions over negative contributions. This makes it particularly robust and suitable for the lower-layer convolutions. Other propagation rules such as LRP- ϵ or LRP-0 are suitable for other layers [126]. Additional propagation rules have been proposed for special layers such as min/max pooling [13], [127], [86] and LSTM blocks [11], [9]. Furthermore, a number of other propagation techniques have been proposed [171], [170], [100], [202] with some of the rules overlapping with LRP for certain choices of parameters. For a technical overview of LRP including a discussion of the various propagation rules and further recent heuristics, see [126].

An inspection of Eq. (2) shows an important property of LRP, that of conserving relevance from layer to layer, in particular, we can show that in absence of bias terms, $\sum_j R_j = \sum_k R_k$. A further interesting property of this propagation rule is 'smoothing': Consider the relevance can be written as $R_j = a_j c_j$ and $R_k = a_k c_k$ a product of activations and factors. Those factors can be directly related by the equation

$$c_j = \sum_k (w_{jk} + \gamma w_{jk}^+) \frac{\max(0, \sum_{0,j} a_j w_{jk})}{\sum_{0,j} a_j (w_{jk} + \gamma w_{jk}^+)} c_k. \quad (3)$$

This equation can be interpreted as a smooth variant of the chain rule for derivatives used for computing the neural network gradient [125]. Thus, analogous to SmoothGrad [176], LRP also performs some gradient smoothing, however, it embeds it tightly into the deep architecture, so that only a

single backward pass is required. In addition to smoothing, Eq. (3) can also be interpreted as a gradient that has been biased to positive values, an idea also found in methods such as *DeconvNet* [201] or *Guided Backprop* [178]. This modified gradient view on LRP can also be leveraged to achieve a simpler and more general implementation of LRP based on ‘forward hooks’, which we describe in the second part of Appendix B, and which we use to apply LRP on VGG-16 [175] and ResNet-50 [65] in Section IV.

E. Other Methods

We discuss in this section several other popular Explainable AI approaches, that either do not fall in the category of post-hoc explanation approaches (and therefore are not covered in the sections above), that are specialized for a particular neural network architecture, or that make use of different units of interpretability than the input features.

In contrast to the discussed post-hoc methods that apply to any DNN model, *self-explainable models* are designed from scratch with interpretability in mind. A self-explainable model can either be trained to solve a machine learning task directly from a supervised dataset, or it can be used to approximate a black-box model on some representative input distribution. Examples of self-explainable models include simple linear models, or specific nonlinear models, e.g. neural networks with an explicit top-level sum-pooling structure [143], [111], [28], [206], [25]. In all of these models, each summand is linked only to one of a few input variables, which makes attribution of their prediction on the input variables straightforward. More complex architectures involving *attention mechanisms* were also proposed [105], [15], [195], and inspection of the attention mechanism itself can also deliver useful insights into the model prediction. While self-explainable models can be useful for many real-world tasks (a list of arguments in favor of these models can be found e.g. in [154]), their applicability becomes more limited when the goal is to explain the strategy of some existing black-box model. In such scenario, one would have to achieve the difficult task of closely replicating the black-box model for every possible input and perturbation of it, while at the same time being constrained by the predefined interpretable structure.

Other methods are *specialized* for a particular deep neural network model for which generic explanation methods do not provide a direct solution. One such model is the graph neural network [160], [91], where the graph adjacency matrix given as input does not appear as it is usually the case in the first layer, but instead at every layer. Methods that have been proposed to explain these particular neural networks include the GNNExplainer [198], or GNN-LRP [162]. Other neural network architectures have a more conventional structure, but still require a non-trivial adaptation of existing explanation methods, for example, extensions of LRP where proposed to deal with the special LSTM blocks in recurrent neural networks [11], [9] or to handle attention units in the context of neural machine translation [38].

Further methods do not seek to explain in terms of input features, but in terms of the latent space, where the directions

in the latent space code for higher-level concepts, such as color, material, object part, or object [205], [17], [18]. In particular, the TCAV method [89] produces a latent-space explanation for every individual prediction. Some techniques integrate multiple levels of abstraction (e.g. different layers of the neural networks), to arrive at more informative explanation of the prediction process [203], [173], [162]. Finally, generative approaches have been proposed to build structured textual explanations of a machine learning model [67], [113].

IV. COMPARING EXPLANATION METHODS

The methods presented in Section III highlight the variety of approaches available for attributing the prediction of a deep neural network to its input features. This variety of techniques also translates into a variety of qualities of explanations. Illustrative examples of images and the explanation of predicted evidence for the ground truth class as produced by the different explanation methods are shown in Fig. 5. Occlusion Analysis is performed by occluding patches of size 32×32 pixels with stride 16. Integrated Gradients performs 5 integration steps starting from 5 random points near the origin in order to add smoothing (cf. Appendix A), resulting in 25 function evaluations. LRP explanations are obtained by applying the same LRP rules as in [126]. We observe the following qualitative properties of the explanations: Occlusion-based explanations are coarse and are indicating relevant regions rather than the relevant pixel features. Integrated Gradients produces very fine pixel-wise explanations containing both substantial amounts of evidence in favor and against the prediction (red and blue pixels). LRP preserves the fine explanation structure but tends to produce less negative scores and attributes relevance to whole features rather than individual pixels.




Fig. 5. Examples of images from ImageNet [157] with classes ‘space bar’, ‘beacon/lighthouse’, ‘snow mobile’, ‘viaduct’, ‘greater swiss mountain dog’. Images are correctly predicted by the VGG-16 [175] neural network, and shown along with an explanation of the predictions. Different explanation methods lead to different qualities of explanation.

In practice, it is important to reach an objective assessment of how good an explanation is. Unfortunately, evaluating explanations is made difficult by the fact that it is generally

impossible to collect ‘ground truth’ explanations. Building such ground truth explanations would indeed require the expert to understand how the deep neural network decides.

Standard machine learning models are usually evaluated by the utility (expected risk) of their decision behavior (e.g. [190]). Transposing this concept of maximizing utility to the domain of explanation, quantifying the utility of the explanation would first require to define what is the ultimate target task (the explanation being only an intermediate step), and then assessing by how much the use of explanation by the human increases its performance on the target task, compared to not using it (see e.g. [14], [62], [41], [159]). Because such end-to-end evaluation schemes are hard to set up in practice, general desiderata for ML explanations have been proposed [184], [123]. Common ones include (1) faithfulness/sufficiency (2) human-interpretability, and (3) possibility to practically apply it to an ML model or an ML task (e.g. algorithmic efficiency of the explanation algorithm).

A. Faithfulness/Sufficiency

A first desiderata of an explanation is to reliably and comprehensively represent the decision structure of the analyzed ML model. A practical technique to assess such property of the model is ‘pixel-flipping’ [158]. The pixel-flipping procedure tests whether removing the features highlighted by the explanation (as most relevant) leads to a strong decay of the network prediction abilities. The procedure is summarized in Algorithm 3.

Algorithm 3 Pixel-Flipping

```

pfcurve = []
for p in argsort(-R) do
     $x \leftarrow x - \{x_p\}$  (remove pixel  $p$  from the image).
    pfcurve.append( $f(x)$ ).
end for
return pfcurve

```

Pixel-flipping runs from the most to the least relevant input features, iteratively removing them and monitoring the evolution of the neural network output. The series of recorded decaying prediction scores can be plotted as a curve. The faster the curve decreases, the more faithful the explanation method is w.r.t. the decision of the neural network. The pixel-flipping curve can be computed for a single example, or averaged over a whole dataset in order to get a global estimate of the faithfulness of an explanation algorithm under study.

Fig. 6 applies pixel-flipping to the three considered explanation methods and on two models: VGG-16 [175] and ResNet-50 [65]. At each step of pixel-flipping, removed pixels are imputed using a simple inpainting algorithm, which avoids introducing visual artefacts in the image.

We observe that for all explanation methods, removing relevant features quickly destroys class evidence. In particular, they perform much better than a random explanation baseline. Fine differences can however be observed between the methods: For example, LRP performs better on VGG-16 than on ResNet-50. This can be explained by VGG-16 having a more

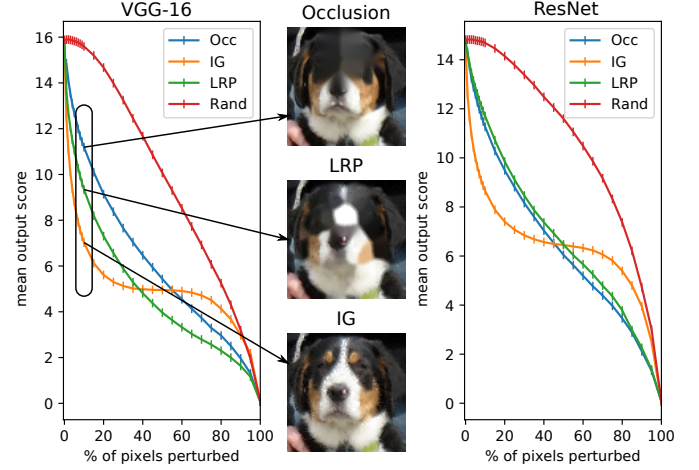


Fig. 6. Pixel-flipping experiment for testing faithfulness of the explanation. We remove pixels found to be the most relevant by each explanation method and verify how quickly the output of the network decreases.

explicit structure (standard pooling operations for VGG-16 vs. strided convolution for ResNet-50), which better supports the process of relevance propagation (see also [149] for a discussion of the effect of structure on the performance of explanation methods).

A second observation in Fig. 6 is that Integrated Gradients has by far the highest decay rate initially but stagnates in the later phase of the pixel-flipping procedure. The reason for this effect is that IG focuses on pixels to which the network is the most sensitive, without however being able to identify fully comprehensively the relevant pattern in the image. This effect is illustrated in Fig. 6 (middle) on a zoomed-in exemplary image of class ‘greater swiss mountain dog’, where the image after 10% flipping has lost most of its prediction score, but visually appears almost intact. Effectively, IG has built an adversarial example [185], [135], i.e. an example whose visual content clearly disagrees with the prediction at the output of the network. We note that Occlusion and LRP do not run into such adversarial examples. For these methods, pixel-flipping steadily and comprehensively removes features until class evidence has totally disappeared.

Overall, the pixel-flipping algorithm characterizes various aspects of the faithfulness of an explanation method. We note however that faithfulness of an explanation does not tell us how easy it will be for a human to make sense of that explanation. We address this other key requirement of an explanation in the following section.

B. Human Interpretability

Here, we discuss whether the presented explanation techniques deliver results that are meaningful to the human, i.e. whether the human can gain understanding into the classifier’s decision strategy from the explanation. Human interpretability is hard to define in general [123]. Different users may have different capabilities at reading explanations and at making sense of the features that support them [147], [133]. For example, the layman may wish for a visual interpretation, even approximate, whereas the expert may prefer an explanation

supported by a larger vocabulary, including precise scientific or technical terms [14].

For the image classification setting, interpretability can be quantified in terms of the amount of information contained in the heatmap (e.g. as measured by the file size). An explanation with a small associated file size is more likely to be interpretable by a human. The table below shows average file sizes (in bytes¹) associated to the various explanation techniques and for two neural networks.

	Occ	IG	LRP
VGG-16	698.4	5795.0	1828.3
ResNet-50	693.6	5978.0	2928.2

We observe that occlusion produces the lowest file size and is therefore the most ‘interpretable’. It indeed only presents to the user rough localization information without going into the details of which exact feature has supported the decision as done e.g. by LRP. On the other side of the interpretability spectrum we find Integrated Gradients. In the explanations this last method produces, every single pixel contains information, and this makes it clearly overwhelming to the human.

In practice, neural networks do not need to be explained in terms of input features. For example, the TCAV method [89] considers directional derivatives in the space of activations (where the directions correspond to higher-level human-interpretable concepts) in place of the input gradient. Similar higher-level interpretations are also possible using the Occlusion and LRP methods, respectively by perturbing groups of activations corresponding at a given layer to a certain concept, or by stopping the LRP procedure at the same layer and pooling scores on some group of neurons representing the desired concept.

C. Applicability and Runtime

Faithfulness and interpretability do not fully characterize the overall usefulness of an explanation method. To characterize usefulness, we also need to determine whether the explanation method is applicable to a range of models that is sufficient large to include the neural network model of interest, and whether explanations can be obtained quickly with finite compute resources.

Occlusion-based explanations are the easiest to implement. These explanations can be obtained for any neural network even those that are not differentiable. This also includes networks for which we do not have the source code and where we can only access their prediction through some online server. Technically, occlusion can therefore be used to understand the predictions of third-party models such as <https://cloud.google.com/vision/> and <https://www.clarifai.com/models>. *Integrated gradients* requires instead for each prediction an access to the neural network gradient. Given that most machine learning models are differentiable, this method is widely applicable also for neural networks with complex structures, such as ResNets [65] or SqueezeNets [79]. Integrated Gradients is also

easily implemented in state-of-the-art ML frameworks such as PyTorch or TensorFlow, where we can make use of automatic differentiation. *Layer-wise Relevance Propagation* assumes that the model is structured as (or can be converted to [85], [86]) a neural network with a canonical sequence of layers, for example, an alternation of linear/convolution layers, ReLU layers, and pooling layers. This stronger requirement and the implementation overhead caused by explicitly accessing the different layers (cf. Appendix B) will however be offset by a last characteristic we consider in this section, which is the computational cost associated producing the explanation. A runtime comparison² of the three explanation methods studied here is given in the table below (measured in *explanations per second*).

	Occ	IG	LRP
VGG-16	2.4	5.8	204.1
ResNet-50	4.0	8.7	188.7

Occlusion is the slowest method as it requires to reevaluate the function for each occluded patch. For image data, the runtime of Occlusion increases quadratically with the step size, making the obtainment of high-resolution explanations with this method computationally prohibitive. Integrated Gradients inherits pixel-wise resolution from the gradient computation which is $O(1)$ but requires multiple iterations for the integration. The runtime is further increased if performing an additional loop of smoothing. LRP is the fastest method in our benchmark by an order of magnitude. The LRP runtime is only approximately three times higher than that of computing a single forward pass. This makes LRP particularly convenient for the large-scale analyses we introduce in Section VI-D where an explanation needs to be produced for every single example in the dataset.

V. THEORETICAL FOUNDATIONS OF EXPLANATION METHODS

In parallel to developing explanation methods that address application requirements such as faithfulness, interpretability, usability and runtime, some works have focused on building theoretical foundations for the problem of explanation [127], [116] and establishing theoretical connections between the different methods [171], [4], [128].

Here, we present three frameworks: the Shapley Values [169], [179], [116] which comes from game theory, the Taylor expansions [13], [19], and the Deep Taylor Decomposition [127], which applies Taylor expansions repeatedly at each layer of a DNN. We then show how Occlusion, Integrated Gradients, or LRP intersect for certain choices of parameters with these mathematical approaches.

A. Shapley Values

Shapley values [169] is a framework originally proposed in the context of game theory to determine individual contributions of a set of cooperating players \mathcal{P} . The method considers

¹JPEG compression using the Pillow image processing library for python with a quality setting of 75/100 (standard settings).

²Explanations are computed in batches of (up to) 16 samples on a GPU and with explanation techniques implemented in PyTorch. Results are averaged over 10 repetitions.

every subset of cooperating players $\mathcal{S} \subseteq \mathcal{P}$ and tests the effect of removing/adding the player i to \mathcal{S} on the total payoff $v(\mathcal{S})$ obtained by \mathcal{S} if they cooperate. Specifically, Shapley values identify the contribution of player p to the overall coalition \mathcal{P} to be:

$$\phi_i = \sum_{\mathcal{S} \subseteq \mathcal{P} \setminus \{i\}} \alpha_{\mathcal{S}} \cdot (v(\mathcal{S} \cup \{i\}) - v(\mathcal{S}))$$

where each subset \mathcal{S} is weighted by the factor $\alpha_{\mathcal{S}} = |\mathcal{S}|! \cdot (|\mathcal{P}| - 1 - |\mathcal{S}|)! / |\mathcal{P}|!$. Shapley values satisfy a number of axioms, in particular, efficiency ($\sum_i \phi_i = v(\mathcal{P})$), symmetry, linearity, and zero added value of a dummy player. Shapley values are in fact the unique assignment strategy that jointly satisfies these axioms [169].

When transposing the method to the task of explaining a machine learning model [179], [116], the players of the cooperating game become the input features, and the payoff function becomes related to the DNN output. In [116], the payoff function is chosen to be the conditional expectation: $v(\mathcal{S}) = \mathbb{E}[f(\mathbf{x}) | \mathbf{x}_{\mathcal{S}}]$. Alternately, to make the score depend only on the model without assuming a specific input distribution, the payoff function can be set to

$$v(\mathcal{S}) = f(\mathbf{x}_{\mathcal{S}}), \quad (4)$$

i.e. input features not in \mathcal{S} are set to zero (see e.g. [5]), and we will use this formulation to make connections to practical explanation methods in Section V-D. Note that Shapley values make almost no assumptions about the structure of the function f and can therefore serve as a general theoretical framework to analyze explanation methods. For specific functions, e.g. additive models of the type $f(\mathbf{x}) = \sum_i f_i(x_i)$, Shapley values (using Eq. (4)) take the simple form $\phi_i = f_i(x_i)$.

B. Taylor Decomposition

Taylor expansions are a well-known mathematical framework to decompose a function into a series of terms associated to different degrees and combinations of input variables. Unlike Shapley values which evaluates the function $f(\mathbf{x})$ multiple times, the Taylor expansion framework for explaining a ML model [13], [19], [42] evaluates the function once at some reference point $\tilde{\mathbf{x}}$ and assigns feature contributions by locally extracting the gradient (and higher-order derivatives). Specifically, the Taylor expansion of some smooth and differentiable function $f: \mathbb{R}^d \rightarrow \mathbb{R}$ at some reference point $\tilde{\mathbf{x}}$ is given by:

$$\begin{aligned} f(\mathbf{x}) &= f(\tilde{\mathbf{x}}) \\ &+ \sum_i [\nabla f(\tilde{\mathbf{x}})]_i \cdot (x_i - \tilde{x}_i) \\ &+ \frac{1}{2} \sum_{i,i'} [\nabla^2 f(\tilde{\mathbf{x}})]_{ii'} (x_i - \tilde{x}_i)(x_{i'} - \tilde{x}_{i'}) \\ &+ \dots \end{aligned}$$

where ∇f and $\nabla^2 f$ denote the gradient and the Hessian respectively, and \dots denote the (non-expanded) higher-order terms. The zero-order term is the function value at the reference point and is zero if choosing a root point. There are as many first-order terms as there are dimensions and each of them is bound to a particular input variable. Thus, they offer a natural way of attributing a function value $f(\mathbf{x})$

onto individual linear components. There are as many second-order terms as there are pairs of ordered variables, and even more third-order and higher-order terms. When the function is approximately locally linear, second and higher-order terms can be ignored, and we get the following simple attribution scheme:

$$R_i = [\nabla f(\tilde{\mathbf{x}})]_i \cdot (x_i - \tilde{x}_i),$$

a product of the gradient and the input relative to our root point. In the general case, there are no closed-form approach to find the root point and it is instead obtained using an optimization technique.

C. Deep Taylor Decomposition

An alternate way of formalizing the problem of attribution of a function onto input features is offered by the recent framework of Deep Taylor Decomposition [127]. Deep Taylor Decomposition assumes the function is structured as a deep neural network and seeks to attribute the prediction onto input features by performing a Taylor decomposition at every neuron of each layer instead of directly on the whole neural network function. Deep Taylor decomposition assumes the output score has already been attributed onto some layer of activations $(a_k)_k$ and attribution scores are denoted by R_k . Deep Taylor Decomposition then considers the function $R_k(\mathbf{a})$ where $\mathbf{a} = (a_j)_j$ is the collection of neuron activations in the layer below. These quantities are illustrated in Fig. 7.




Fig. 7. Graphical illustration of the function $R_k(\mathbf{a})$ that DTD seeks to decompose on the input dimensions. Because R_k is complex, it is often replaced by an analytically more tractable model $\hat{R}_k(\mathbf{a})$ that only depends on local activations.

The function $R_k(\mathbf{a})$ is typically very complex as it corresponds to a composition of multiple forward and backward computations. This function can however be approximated locally by some ‘relevance model’ $\hat{R}_k(\mathbf{a})$, the choice of which will depend on the method we have used for computing R_k . We then compute a Taylor expansion of this function:

$$\begin{aligned} \hat{R}_k(\mathbf{a}) &= \hat{R}_k(\tilde{\mathbf{a}}) \\ &+ \sum_j [\nabla \hat{R}_k(\tilde{\mathbf{a}})]_j \cdot (a_j - \tilde{a}_j) \\ &+ \dots \end{aligned}$$

The linear terms define ‘messages’ $R_{j \leftarrow k}$ that can be redistributed to neurons in the lower layer, and messages received by a given neuron at a certain layer are summed to form a total relevance score:

$$R_j = \sum_k [\nabla \hat{R}_k(\tilde{\mathbf{a}}^{(k)})]_j \cdot (a_j - \tilde{a}_j^{(k)}) \quad (5)$$

here, we have added an index $\{ \}^{(k)}$ to the root point to make explicit that different root points can be used for expanding different neurons. The redistribution procedure is iterated from the top layer towards the lower layers, until the input features are reached.

D. Connections between Explanation Methods

Having described the Shapley values and the simple and deep Taylor decomposition frameworks, we now present some results from the literature showing how some explanation methods reduce for certain choices of parameters to these frameworks. The different connections we outline here are summarized in Table I.

TABLE I
EXAMPLES OF EXPLANATION METHODS APPLIED WITH DIFFERENT PARAMETERS ON DIFFERENT MODELS, AND WHETHER THEY CAN BE EMBEDDED IN EACH OF THE THREE PRESENTED THEORETICAL FRAMEWORKS.

	Shapley	Taylor	DTD
<i>Linear models</i>			
Occlude-1	✓	✓	(✓)
IG on Segment($\mathbf{0}, \mathbf{x}$)	✓	✓	(✓)
LRP-0	✓	✓	(✓)
<i>Nonlinear additive models</i>			
Occlude-1	✓		
<i>Deep rectifier networks</i>			
Occlude-1			
IG on Segment($\mathbf{0}, \mathbf{x}$)		✓	
LRP-0		✓	✓
LRP- ϵ/γ ...			✓

We start by connecting occlusion-based explanations of a linear model to Shapley values and Taylor decomposition.

Proposition 1. *When applied to homogeneous linear models (of the type $f(\mathbf{x}) = \mathbf{w}^\top \mathbf{x}$), occlusion with patch size 1 and replacement value 0 is equivalent to a Taylor decomposition with root point $\tilde{\mathbf{x}} = \mathbf{0}$, as well as Shapley values with value function given by Eq. (4).*

The first connection is shown by the chain of equations $f(\mathbf{x}) - f(\mathbf{x} - \{x_i\}) = w_i x_i = [\nabla f(\mathbf{0})]_i \cdot (x_i - 0)$. For the Shapley values, we simply observe $\phi_i = \sum_{S \subseteq \mathcal{P} \setminus \{i\}} \alpha_S \cdot (w_i \phi_i) = 1 \cdot (w_i \phi_i)$, which again gives the same result. Integrated gradients and LRP-0 also yield the same result. Hence, for this simple linear model, all explanation methods behave consistently and in agreement with the existing theoretical frameworks. The connection of explanation methods to Shapley values for linear models was also made in [5].

The connection between Integrated Gradients and Taylor decomposition holds for a broader class of neural network functions, specifically deep rectifier networks (without biases):

Proposition 2. *When applied to deep rectifier networks of the type $f(\mathbf{x}) = \rho(W_L \rho(\dots \rho(W_2 \rho(W_1 \mathbf{x}))))$, Integrated Gradients with integration path $\{t\mathbf{x}; 0 < t \leq 1\}$ is equivalent to Taylor decomposition at $\tilde{\mathbf{x}} = \varepsilon \mathbf{x}$ in the limit $\varepsilon \rightarrow 0$.*

This can be shown by making the preliminary observation that a deep rectifier network is linear with constant gradient on

the segment $(\mathbf{0}, \mathbf{x}]$ and then applying the chain of equations $\int_\varepsilon^1 x_i [\nabla f(t\mathbf{x})]_i dt = (1 - \varepsilon) x_i [\nabla f(\varepsilon \mathbf{x})]_i = [\nabla f(\varepsilon \mathbf{x})]_i (x_i - \varepsilon x_i)$. This connection, along with the observation that a single gradient evaluation of a deep network can be noisy (cf. Section II-B) speaks against integrating on the segment $(\mathbf{0}, \mathbf{x}]$. For this reason, we have opted in the experiments of Section IV to use a smoothed version of IG. A further result shows an equivalence between a ‘naive’ version of LRP (using LRP-0 at every layer) and Taylor decomposition.

Proposition 3. *For deep rectifier nets of the type $f(\mathbf{x}) = \rho(W_L \rho(\dots \rho(W_2 \rho(W_1 \mathbf{x}))))$, applying LRP-0 at each layer is equivalent to a Taylor decomposition at $\tilde{\mathbf{x}} = \varepsilon \mathbf{x}$ in the limit $\varepsilon \rightarrow 0$.*

This result can be derived by taking the LRP formulation of Eq. (3) and setting $\gamma = 0$. This equation then reduces to:

$$c_j = \sum_k w_{jk} \text{step}(\sum_{0,j} a_j w_{jk}) c_k$$

where $\text{step}(t) = 1_{t>0}$. This equation is exactly the same as the one that propagates gradients in a deep rectifier network. Hence, the input relevance computed by LRP becomes $R_i = x_i c_i = x_i [\nabla f(\mathbf{x})]_i$ for which we have already shown the equivalence to simple Taylor decomposition in the proposition above.

Proposition 4. *For deep rectifier networks of the type $f(\mathbf{x}) = \rho(W_L \rho(\dots \rho(W_2 \rho(W_1 \mathbf{x}))))$, applying LRP- γ is equivalent to performing one step of deep Taylor decomposition and choosing the nearest root point on the line $\{\mathbf{a} - t\mathbf{a} \odot (1 + \gamma \cdot \mathbf{1}_{\mathbf{w}_k \geq 0}); t \in \mathbb{R}\}$.*

We choose the relevance model $\hat{R}_k(\mathbf{a}) = a_k(\mathbf{a}) \cdot c_k$ with c_k constant (cf. [126] for a justification). Injecting the root point in the first-order terms of DTD (summands of Eq. (5)) gives:

$$\begin{aligned} R_{j \leftarrow k} &= w_{jk} \cdot c_k \cdot (a_j - (a_j - t a_j \cdot (1 + \gamma \cdot \mathbf{1}_{w_{jk} \geq 0}))) \\ &= a_j \cdot (w_{jk} + \gamma w_{jk}^+) \cdot t \cdot c_k \end{aligned}$$

where t is resolved using the conservation equation $\sum_j R_{j \leftarrow k} = R_k$. LRP-0 is a special case of LRP- γ with $\gamma = 0$. A similar procedure with another choice of reference point gives LRP- ϵ (cf. [126]).

VI. EXTENDING EXPLANATIONS

The explanation methods we have presented in the previous sections were applied to a particular class of models (deep neural networks) and produced particular types of explanations (attribution on the input features). We present in the following various extensions that broaden the applicability of these methods and diversify the type of explanation that can be produced. In particular, we will discuss (1) higher-order methods to produce richer explanations involving combination of features, (2) a systematic way of extending explanation methods to non-neural network models e.g. in unsupervised learning where explanations are also needed, (3) a principled way to ensure that explanations of DNN classifiers are class-discriminative, and (4) strategies to go beyond individual explanations to arrive at a general understanding of the ML model.

A. Explaining Beyond Heatmaps

The locally linear structure of deep neural networks lends itself well to the heatmap-based methods we have reviewed in this paper, which we call first-order methods. However, special types of neural networks, e.g. that incorporate products between input or latent variables lose that property. Neural networks with product structures commonly occur for relational tasks such as comparing images [122] or collaborative filtering [66]. Graph neural networks [160], [91] multiply the input connectivity matrix multiple times, and consequently also exhibit product structures. In that case, the neural network is no longer piecewise linear and typically becomes piecewise polynomial with its input.

For illustration, we present the BiLRP method [42], which assumes we have a similarity model built as a dot product on some hidden representation $\phi = \phi_L \circ \dots \circ \phi_1$ of a deep network:

$$y(\mathbf{x}, \mathbf{x}') = \langle \phi_L \circ \dots \circ \phi_1(\mathbf{x}), \phi_L \circ \dots \circ \phi_1(\mathbf{x}') \rangle$$

If all functions ϕ_1, \dots, ϕ_L are piecewise homogeneous linear, then y can be rewritten as a composition:

$$y(\mathbf{x}, \mathbf{x}') = \psi_L \circ \dots \circ \psi_1(\mathbf{x}, \mathbf{x}')$$

with ψ_1, \dots, ψ_L piecewise bilinear. Using deep Taylor decomposition, but at each step applying a *second-order* Taylor expansion, we arrive conceptually at an attribution of the similarity score $y(\mathbf{x}, \mathbf{x}')$ to *pairs* of input features. (Practically, the attribution can be expressed as a product of two branches of LRP computation, hence the name BiLRP.) An example of BiLRP explanation for the similarity of two planes in VGG-16 feature space is shown in Fig. 8.



Fig. 8. Example of two images predicted to be similar, along with a BiLRP second-order attribution of their similarity score rendered as a bipartite graph. (Figure from [42]). The explanation shows that the front part of the two planes jointly contributes to the predicted similarity.

We observe that the explanation does not highlight individual features but instead pairs of features from the two input images, reflecting the fact that they are jointly relevant to explain the high similarity score.

Other higher-order methods include GNN-LRP which explains graph neural networks in terms of collections of edges [162], or Integrated Hessians [81], which can be seen as a second-order extension of Integrated Gradients.

Finally, we note that the Shapley framework also offers the possibility to quantify the joint contribution of two interacting players through the Shapley interaction index [55]. Lundberg et al. [115] recently applied this concept to compute higher-order explanations of a machine learning model. Here the

interaction value between feature i and feature j can be computed as

$$\phi_{ij} = \sum_{S \subseteq \mathcal{P} \setminus \{i, j\}} \alpha_S \cdot (v(S \cup \{i, j\}) - v(S \cup \{i\}) - v(S \cup \{j\}) + v(S))$$

with the weight factor $\alpha_S = |\mathcal{S}|! \cdot (|\mathcal{P}| - 2 - |\mathcal{S}|)! / 2^{|\mathcal{P}| - 1}!$. Shapley interaction values allow for a separate consideration of interaction effects, which can be crucial for understanding the model prediction strategy (see Section VII-C for a worked-through example in a medical application). However, although theoretically valid, these Shapley interaction terms are in most cases infeasible to compute for complex machine learning models such as deep neural networks, and approximations are therefore required [120].

B. Explaining Beyond Deep Networks

Deep neural networks have been shown to perform extremely well on classification or regression tasks. However for unsupervised problems such as anomaly detection or clustering, although deep models have also reached successes [27], [156], [155], shallow models (e.g. centroid-based [119], PCA-based [64], kernel-based [141], or combinations of them [164], [37], [72]) remain popular workhorses. As these models are not given in the form of a neural network, a direct application of methods designed in the context of linear models and DNNs is not feasible.

While a possible approach to explaining would be to train a surrogate neural network to match the output of the unsupervised model, it was found that some unsupervised models such as k-means clustering [119] or kernel density estimation [141] can be directly rewritten as a neural network (or ‘neuralized’), without requiring any retraining or architecture search [85], [86].

Consider for illustration a kernel k-means model of the type studied in [37]. For this type of model, and assuming a Gaussian kernel $\mathbb{K}(\mathbf{x}, \mathbf{x}') = \exp(-\gamma \|\mathbf{x} - \mathbf{x}'\|^2)$, the probability ratio in favor a given cluster ω_c can be expressed as:

$$\frac{P(\omega_c | \mathbf{x})}{1 - P(\omega_c | \mathbf{x})} = \frac{(Z_c^{-1} \sum_{i \in \mathcal{C}_c} \mathbb{K}(\mathbf{x}, \mathbf{x}_i))^{\beta/\gamma}}{\sum_{k \neq c} (Z_k^{-1} \sum_{j \in \mathcal{C}_k} \mathbb{K}(\mathbf{x}, \mathbf{x}_j))^{\beta/\gamma}} \quad (6)$$

This is a power-assignment model applied to the kernel density functions of each cluster. The sets \mathcal{C}_c and \mathcal{C}_k are the representatives for clusters c and k , and Z_c, Z_k are respective normalization factors. An example of decision function produced by this model for a three-cluster problem is shown in Fig. 9 (left). Clearly, Eq. (6) is a priori not composed of neurons. However, it can be reorganized into the following sequence of detection and pooling functions [85]:

$$\log \left[\frac{P(\omega_c | \mathbf{x})}{1 - P(\omega_c | \mathbf{x})} \right] = \beta \min_{k \neq c}^\beta \left\{ \min_{j \in \mathcal{C}_k}^\gamma \left\{ \max_{i \in \mathcal{C}_c}^\gamma \{ \mathbf{w}_{ij}^\top \mathbf{x} + b_{ijk} \} \right\} \right\}$$

with $\mathbf{w}_{ij} = 2(\mathbf{x}_i - \mathbf{x}_j)$ and $b_{ijk} = \|\mathbf{x}_j\|^2 - \|\mathbf{x}_i\|^2 + \gamma^{-1}(\log Z_k - \log Z_c)$ are parameters of the first linear layer. This layer is followed by a hierarchy of log-sum-exp computations interpretable as canonical max- and min-pooling operations. The neuralized version of kernel k-means is depicted in Fig. 9 (right).




Fig. 9. *Left*: Kernel k-means applied to a toy two-dimensional problem with three clusters. Red and blue color in the background represent the positive and negative values of the logit function for a given cluster. *Right*: 4-layer neural network equivalent of the kernel k-means logit score [85].

With this structure, explanation techniques such as LRP can be used to produce explanations of cluster membership [85]. In particular, the LRP backward pass first identifies the most relevant class competitors, then the most relevant representatives (i.e. data points) of these class competitors and of the cluster of interest, and finally the most relevant directions in input space, thereby producing a heatmap-based explanation of the cluster assignment [85].

C. Explaining Beyond Output Neurons

The concept of neuralization can also be applied in a supervised learning setting, for improving the explanation of a deep neural network classifier. So far, we have explained quantities at the output of the last linear layer of the network. Because these output quantities are unnormalized they may respond positively to several classes, thereby lacking selectivity. The problem of class selectivity was highlighted e.g. in [57], [80], [126] along with practical solutions to overcome this effect. Here, we present the ‘neuralization’ approach of [126], which first makes the observation that ratios of probabilities as given by the top-layer soft-assignment model can be expressed as:

$$\frac{P(\omega_c|\mathbf{x})}{1 - P(\omega_c|\mathbf{x})} = \frac{\exp(\mathbf{w}_c^\top \mathbf{a})}{\sum_{k \neq c} \exp(\mathbf{w}_k^\top \mathbf{a})}$$

This computation can then be reorganized in the two-layer neural network

$$\log \left[\frac{P(\omega_c|\mathbf{x})}{1 - P(\omega_c|\mathbf{x})} \right] = \min_{k \neq c} \{(\mathbf{w}_c - \mathbf{w}_k)^\top \mathbf{a}\}$$

where min is a soft minimum implemented by a log-sum-exp computation. The DNN processing up to the output neuron or up to the output of the neuralized logit model is illustrated in Fig. 10 along with LRP explanations for these two quantities associated to the class ‘passenger_car’.

In the first explanation, both the passenger car and the locomotive can be seen to contribute. In the second explanation, the locomotive turns blue. The latter is indeed speaking for the class locomotive, which mechanistically lowers the probability for the class ‘passenger_car’ [126]. This example shows that it is important in presence of correlated features to precisely define what quantity (unnormalized score or logit) we would like to explain.

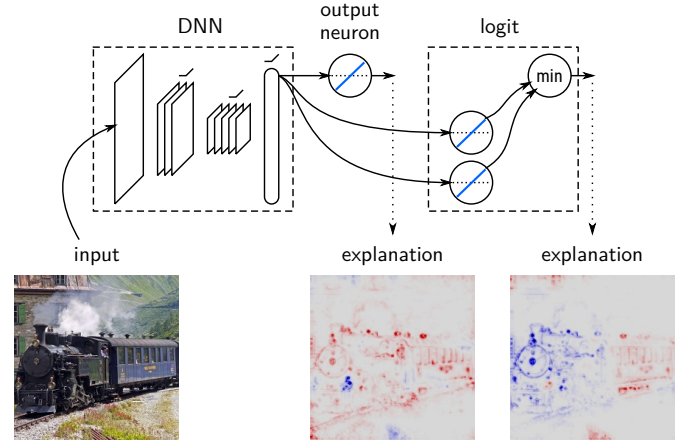


Fig. 10. Deep neural network to which we append a neuralized version of the log-likelihood ratio [126]. Considering the latter quantity instead of the DNN output leads to a different explanation.

We note that while neuralization has served here to support LRP-type explanations, the concept could potentially be used for other purposes. The identified neural network structure may help to gain further understanding of the model or provide intermediate representations that are potentially useful to solve related tasks.

D. Explaining Beyond Individual Predictions

In practice, we may not only be interested in explaining how the DNN predicts a single data point, but also in the statistics of them for a whole dataset. This may be useful to validate the model in a more complete manner. Let $f: \mathbb{R}^d \rightarrow \mathbb{R}$ be a function that takes a data point as input and predicts evidence for a certain class for each data point. Consider a dataset $\mathbf{x}_1, \dots, \mathbf{x}_N$ of such data points. The total class evidence can be represented as a function $g: \mathbb{R}^{N \times d} \rightarrow \mathbb{R}$ where:

$$g(\mathbf{x}_1, \dots, \mathbf{x}_N) = \sum_{n=1}^N f(\mathbf{x}_n)$$

This composition of the neural network output and a sum-pooling remains explainable by all methods surveyed here, however, the explanation is now high-dimensional ($N \times d$).

1) *Relevance Pooling*: Practically, we may not be interested in explaining every single data point in terms of every single input features. A more relevant information to the user would be the overall contribution of a subgroup of features \mathcal{I} on a group of data points \mathcal{G} (cf. [101], [128]). In particular the Integrated Gradient and LRP methods surveyed here produce explanations that satisfy the conservation property:

$$g(\mathbf{x}_1, \dots, \mathbf{x}_N) \approx \sum_{n=1}^N \sum_{i=1}^d R_{i,n}$$

and that can be converted to a coarse-grained explanation

$$\approx \sum_{\mathcal{G}} \sum_{\mathcal{I}} \underbrace{\sum_{n \in \mathcal{G}} \sum_{i \in \mathcal{I}} R_{i,n}}_{R_{\mathcal{I}, \mathcal{G}}}$$

that still satisfies the desired conservation property. As an illustration of the concept, we consider the ‘Concrete Compressive Strength’ example of Section II. Data points are grouped in three k-means clusters, and features are grouped in two sets: the singleton {age}, and the set of all remaining features describing concrete composition. The pooled analysis is illustrated in Fig. 11.

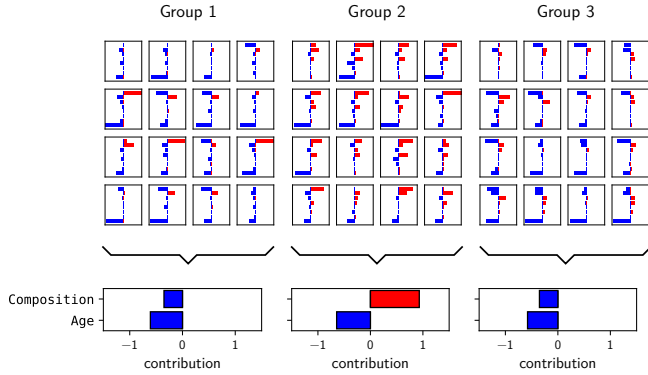


Fig. 11. Pooled analysis. *Top*: Feature-wise contributions for the prediction on three clusters of the Concrete Compressive Strength Dataset [197]. *Bottom*: Coarse-grained explanations obtained by pooling contributions on data clusters and groups of features.

This analysis gives further insight into our predictive model. We observe that most distinguishing factors, especially age, contribute negatively to strength. In other words, a ‘typical’ age and composition is a recipe for strength whereas high/low values tend to be explanatory for weakness. Notably, one data cluster stands out by having composition features that are explanatory for strength.

2) *Spectral Relevance Analysis (SpRAY)* [104]: While in Section VI-D1 we have reduced the dimensionality through pooling, other analyses are possible. For example, the SpRAY method [104] does not assume a fixed pooling structure (e.g. a partition of data points and a partition of features), and applies instead a clustering of explanations in order to identify prototypical decision strategies. Algorithm 4 outlines the three steps procedure used by SpRAY:

Algorithm 4 Spectral Relevance Analysis

```

for  $n = 1$  to  $N$  do
   $\mathbf{R}^{(n)} \leftarrow \text{explain}(\mathbf{x}^{(n)}, f)$ 
   $\bar{\mathbf{R}}^{(n)} \leftarrow \text{normalize}(\mathbf{R}^{(n)})$ 
end for
clustering( $\{\bar{\mathbf{R}}^{(1)}, \dots, \bar{\mathbf{R}}^{(N)}\}$ )

```

The method first produces an explanation for each data point. In principle, any explanation method can be used, e.g. occlusion, integrated gradients, or LRP. Explanations are then normalized (e.g. blurred and standardized) to become invariant to small pixel-wise or saliency variations. Finally, a clustering algorithm is applied to the normalized explanation, and examples with the same cluster index can be understood as being associated with some prototypical decision strategy, e.g. looking at the object, looking at the background, etc.




Fig. 12. SpRAY analysis. Explanation of the predictions of the Concrete Compressive Strength Dataset [197], displayed at coordinates corresponding to their t-SNE embedding. This analysis provides a visualization of the overall classifier’s strategy.

Alternately, the clustering step can be replaced by a low-dimensional embedding step to produce a visual map of the overall decision structure of the ML model. The SpRAY analysis is illustrated in Fig. 12 on the same Concrete Compressive Strength Dataset [197] used before.

Here, we observe, for example, that a typical decision strategy (bottom left) consists of explaining high compressive strength based on the first feature (high cement concentration), whereas another typical decision strategy (bottom right) consists of predicting low compressive strength based on the last feature (low concrete’s age). This rich visual feedback can serve to validate the prediction strategy with expert knowledge to make sure that the model does not predict based on a dataset artefact. We will present in Section VIII-A a successful real-world use-case of the SpRAY analysis for inspecting a state-of-the-art model trained on a large image classification dataset.

Altogether, relevance pooling and SpRAY support a variety of dataset-wide analyses that are useful to explore and characterize the decision strategies of complex models trained on large datasets.

VII. WORKED-THROUGH EXAMPLES

In this paper, we have motivated the use of explanation in the context of deep learning models and showcased some methods for obtaining explanations. Here, we aim to take a practical look for the user to assess when explanation is required, what are common issues with applying explanation techniques / setting their hyperparameters, and finally, how to make sure that the produced explanations deliver meaningful insights for the human.

A. Example 1: Validating a Face Classifier

In the first worked-through example we wish to train an accurate classifier for predicting a person’s age from images of faces. We will show how to use explanation for this task,

in particular, to verify that the model is not using “wrong” features for its decisions.

Let us use for this the Adience benchmark dataset [44] providing 26,580 images captured ‘in the wild’ and labelled into eight ordinal groups of age ranges $\{(0-2), (4-6), (8-13), (15-20), (25-32), (38-43), (48-53), (60+)\}$.

Because the number of examples in this dataset is limited and likely not sufficient to extract good visual features, we adopt the common approach of starting with a generic pre-trained classifier and fine-tune it on our task. We download a VGG-16 [175] neural network architecture pretrained on ImageNet [35] obtainable from `modelzoo.co`. First test results after training using Stochastic Gradient Descent (SGD) [108] report reasonable performance, with *exact* and *1-off* [168], [44] prediction accuracy³ of 56.5% and 90.0%, respectively. Here, the 1-off accuracy considers predictions of (up to) one age group away from the true label as correct.

In order to understand the learned prediction strategies of our model and to verify that it uses meaningful features in the training data, we test different explanation methods, first, a simple occlusion and gradient analysis, and we also take an off-the-shelf explanation software, the LRP Toolbox [102] for Caffe [83], and choose the method LRP configured to perform ‘LRP- ϵ ’ on all layers in a first attempt. Explanations are shown for a given image in Fig. 13 (a).

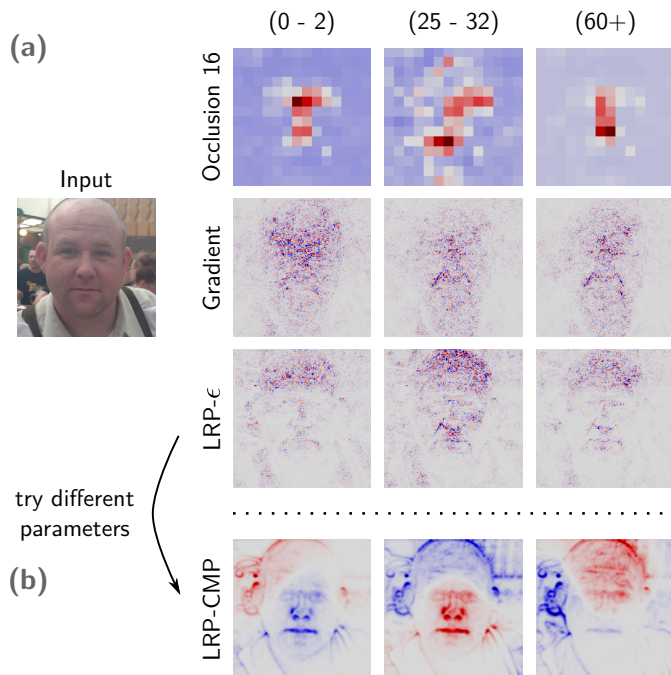


Fig. 13. *Top*: Explanations produced by different methods on the VGG-16 age model. Results are shown for the output neurons associated to age group labels (0-2), (25-32) and (60+) respectively. *Bottom*: Application of the layer-dependent LRP-CMP decomposition strategy.

Some insight can be readily obtained from these explanations, e.g. the classifier has learned to attend the center of the image where the face is. However, we also observe

³Results have been averaged over the official pre-selected five-fold dataset split [44].

several limitations of the explanations: Here, the occlusion-based analysis produces a blue background (negatively relevant) instead of identifying the background as irrelevant. The next two explanations (gradient and LRP- ϵ) are here overly complex with frequent local sign changes, making it hard to extract further insights, especially what are the features that contribute to different age groups. The reason for such poor performance is that we have either not considered a broad enough spectrum of explanation methods, or the parameters of these methods have not been set optimally. This leads to our first recommendation:

Try different parameters of the explanation techniques

Specifically, we will now try an alternate LRP preset called ‘LRP-CMP’ that applies a composite strategy [103], [126], [94] where different rules are applied at different layers. Explanations obtained with this new rule are given in Fig. 13 (*bottom*). The new explanations highlight features in a much more interpretable way and we also start to better understand what speaks — according to the model — in favor of or against certain age groups. For example, explanations amusingly reveal baldness as a feature corresponding to both age groups (0-2) and (60+). In the shown sample, baldness contributes evidence for the classes (0-2) and (60+), while it speaks against the age group (25-32). Relatedly, the expression of the man’s chin and mouth area contradicts class (0-2) more than class (60+), but ‘looks like’ it would belong to a person aged (25-32).

Let’s now move back to the initial question, namely how to verify that the model is using the right features for predicting. While the decision structure of the model was meaningful in Fig. 13, we would like to verify it is also the case for other test cases. Fig. 14 (*top*) shows further examples from the Adience dataset; a woman labelled (60+) and three images of the same male labelled (48-53) with smiles of varying intensities.



Fig. 14. LRP heatmaps demonstrating the effects of ImageNet [35] pretraining (*middle*) compared to additional IMDB-WIKI [153] pretraining (*bottom*). All heatmaps show the model decision w.r.t. age group (60+).

We apply LRP with the same preset ‘LRP-CMP’ on these images. LRP evidence for each image for the class (60+)

is shown in Fig. 14 (middle). Surprisingly, according to the model, broad smiling contradicts the prediction of belonging to the age group (60+). Smiling is however clearly a confounding factor, which reliably predicts age group only to the extent that no such case is present in the training data. This predicting strategy is related to the ‘Clever Hans’⁴ effect [104] and we can therefore formulate our second recommendation:

Unmask ‘Clever Hans’ examples

Note that instead of screening through all images manually, we can also use techniques such as SpRAY [104], which perform such analysis semi-automatically for large datasets such as ImageNet (see also Section VIII-A for successful applications).

While for the examples showcased in Fig. 14 other features may compensate for the ‘smiling’ effect, — here almost all features other than the smile also affect the decision towards this age group positively, — this will cause errors for less clear-cut cases. This may explain why the accuracy of the ImageNet-based model is not very high, and can point at the fact that the test set accuracy may drop dramatically on new datasets, e.g. comprising more old people smiling.

Instead, we would like our model to be robust to a subject’s mood when predicting his or her age. We thus need to find a way to prevent Clever Hans effects, e.g. prevent the model to associate smiling with age. One reason the model has learned that connection in the first place is the extreme population imbalance among the age groups of the Adience dataset; a problem which is shared with many other datasets of face images, e.g. [78], [153]. We therefore add a second pre-training phase in between the ImageNet initialization and the *actual* training based on the Adience data, by using the considerably larger IMDB-WIKI [153] dataset. The IMDB-WIKI dataset consists of 523,051 images from 20,284 celebrities (460,723 images from the Internet Movie Data Base (Imdb) and 62,328 images from Wikipedia) at different ages, labelled with 101 labels (0-100 years, one label per year). The IMDB-WIKI dataset also suffers from highly imbalanced label populations. However, we follow [153] and re-normalize the age distribution by under-sampling the more frequent classes until approximately 260,000 samples are selected overall. Furthermore, we assume that since the IMDB-WIKI dataset is composed of photos of public figures (taken at publicized events) the ratio of expressed smiles in higher age groups will be more frequent than in the Adience dataset, which has been captured ‘in the wild’. A comparison of performance on the Adience benchmark of the original model (pretrained on ImageNet only) and the improved model is given in the table below.

	accuracy	1-off
ImageNet pretrained	56.5	90.0
IMDB-WIKI pretrained	63.0	96.0

We observe that the additional and more domain-specific IMDB-WIKI pretraining step has improved the generalization performance of the VGG-16 model. Furthermore, we will see that it also prevented the model from associating smiling exclusively with younger age groups. Fig. 14 (*bottom*) shows LRP heatmaps for all four examples and age label (60+). For the woman, the model has shifted its attention from the hair and clothes to the face region and neck, and no longer considers the smile as contradictory to the class. A similar effect can be observed for the samples showing the male person. The model’s age prediction capabilities can no longer be fooled by just smiling into the camera. However, by introducing the IMDB-WIKI pretraining step, we have apparently replaced the smile-related Clever Hans strategy with another one, related to the presence of glasses in images of males in higher age groups. This leads to our third recommendation:

Iteratively validate and improve the model

We can do so until the model solely relies for its predictions on meaningful face features. For that, choosing a better pretraining may not be sufficient, and other more advanced interventions may be required.

B. Example 2: Identifying Male and Female Speech Features

After demonstrating how explanations can be used to unmask Clever Hans strategies, or more generally validate a classifier, we will now discuss another use case, where explanations are this time applied not to get a better model, but to gain new (scientific) insights. In this worked-through example, we will show that explanations can be used to identify male and female features in speech.

Before going into the analysis, let us first introduce the data and the model used for the speaker’s male vs. female classification task. As training data we use the recently recorded AudioMNIST [20] dataset, comprised of 30000 audio recordings of spoken digits from 60 different speakers, with 50 repetitions per digit and speaker, in a 48kHz sampling frequency. Next to annotations for spoken digit (0-9) and sex of speaker (48 male, 12 female), the dataset provides labels for speaker age, accent and origin. We begin by training a deep neural network model on the raw waveform data, which is first downsampled to 8kHz, and randomly padded with zeroes before and after the recorded signal to obtain a 8000 dimensional input vector per sample. A CNN architecture comprised of six 1d-convolution layers interleaved with max-pooling layers and topped of with three fully connected layers [20] and ReLU activation units after each weighted layer is prepared for optimization. In order to prevent the model from overfitting on the more frequent population of samples labelled as ‘male’, we (randomly) select 12 speakers from both classes. The model is then trained and evaluated in a 4-fold cross-validation setting, in which the 24 speakers are grouped into four sets of 3 male speakers and 3 female speakers. Each of the four splits thus contains 1000 waveform

⁴‘Clever Hans’ was a famous horse at the beginning of the 20th century, which was believed by his trainer to be capable of performing arithmetic calculations. Subsequent analyses revealed that the horse was not performing arithmetic calculations but was detecting cues on the face of his trainer to produce the right answers. In machine learning, the term ‘Clever Hans’ can be used to designate strategies that mimic the expected behavior but are based on unexpected correlations or artefacts in the data [104].

features. Two folds are used for training, while one of the remaining data splits are reserved for validation and testing. The model reaches an average test set accuracy (\pm standard deviation) of $91.74\% \pm 8.60\%$ across all splits.

With the goal of understanding the data better by explaining the model, we consider two examples predicted by the network to be male and female and apply LRP to visualize those predictions. Here, the waveform is represented as a scatter plot where each time step is color-coded by its relevance. Results are shown in Fig. 15.

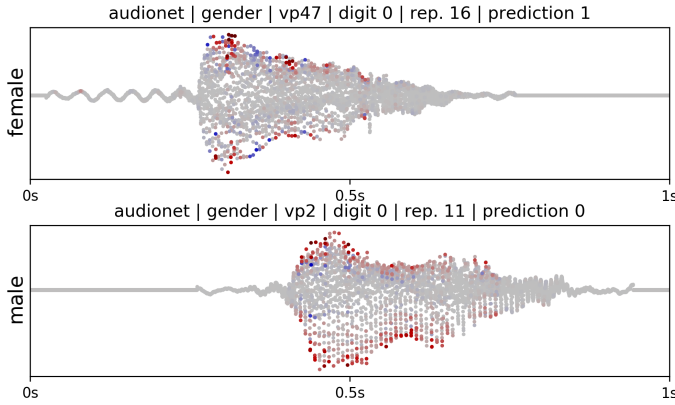


Fig. 15. Explanations based on waveform representation of speech data. Correct prediction of a female (*top*) and male (*bottom*) subject. The waveform data is visualized as a scatter plot of 8000 discrete measurements, color coded according to relevance attribution for the true class label.

The explanations reveal that the model predominantly uses the *outer hull* of the waveform signal for decision making. For a human observer, however, these explanations are difficult to interpret due to the limited accessibility of the data representation in the first place (see Fig. 15). Although the model performs reasonably well on waveform data, it is hard to obtain a deeper understanding beyond the network’s *modus operandi* based on relevance maps, due to the limitations imposed by the data representation itself. We therefore opt to change the data representation for improved interpretability.

Make your input features interpretable

More precisely, we exchange the raw waveform representation of the data with a corresponding 228×230 (time \times frequency) shaped spectrogram representation by applying a short-time Fourier transform (time segment length of 455 samples, with 420 samples overlap), cropped to a 227×227 matrix by discarding the highest frequency bin and the last three time segments. Consequently we also exchange the neural network architecture and use an AlexNet [99] model, which is able to process the transformed input data using 2d-convolution operators.

Fig. 16 visualizes two input spectrograms, with corresponding relevance maps (only relevance values with more than 10% relative amplitude) drawn on top.

Heatmap visualizations based on spectrogram input data are more informative than those for waveform data and reveal that the model has learned to distinguish between male and female speakers based on the lowest fundamental frequencies (male speakers, Fig. 16 (*right*)), and immediate harmonics (female

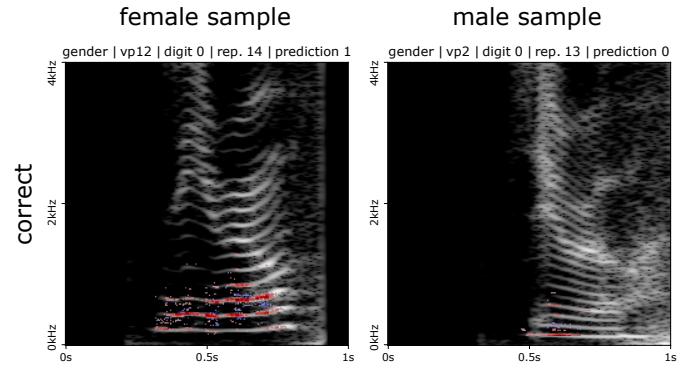


Fig. 16. *Left*: Spectrogram representation of digit ‘zero’ spoken by female speaker ‘vp12’. *Right*: Spectrogram representation of digit ‘zero’ spoken by male speaker ‘vp2’. Relevance maps are shown w.r.t. to the samples’ true classes.

speakers, Fig. 16 (*left*)) shown in the spectrogram. Many incorrectly classified samples with ground truth label ‘male’ show large gaps between frequency bands often occurring in samples from female speakers. Note that these insights are consistent with the literature [187].

Gain insights by explaining predictions

As a noteworthy side effect, the increase in interpretability from switching from a waveform data representation to spectrogram data representation does *not* come at a price of model performance. On the contrary, model performance is even increased slightly from 91.74% to 95.87%.

For applications where a change of input representation into images is not possible, e.g. because the training data are not available or because the human-interpretable image representation does not allow for training an accurate model, we can try to increase interpretability by different means. For our speech example, that could consist of segmenting the time series into meaningful phases, as done in ECG analysis, and compute relevance w.r.t. these phases rather than the individual time points. Furthermore, the data and the explanations can be projected in different domains (e.g. similar to the TCAV approach [89]). In our speech example, projecting the time series and the explanations on a Fourier basis would for example allow for understanding frequencies that really matter for the decision. All this “post-preprocessing” of the explanation can help to understand the results, even though the input features are hardly interpretable.

C. Example 3: Identifying Interacting Mortality Risk Factors

The third worked-through example demonstrates possible limitations of first-order explanation methods and the need for a detailed (higher-order) analysis of explanation results, especially when aiming to draw causal conclusions on the basis of the explanations. The example is adapted from the work of Lundberg et al. [115]. It concerns a medical application, in particular, a statistical model used for predicting mortality called ‘cox proportional hazards model’ [33], fitted on data from NHANES I with follow-up mortality data from the NHANES I Epidemiologic Follow-up Study [32]. The

mortality dataset consists of 14,407 individuals and 79 features and the authors use a Tree-based model⁵ and a Shapley additive explanation (SHAP) method.

Fig. 17 (a) shows the computed SHAP values for one of the considered features, namely the age of the participants. Consistent with common medical knowledge, the SHAP values indicate a higher mortality risk for people with higher age. The vertical dispersion at a single age value results from interaction effects in the model, which are hidden when only considering the first-order explanations. In this dataset the sex of the participant modulates the risk associated with the factor “age”. This interesting relation can be derived from the SHAP interaction values.

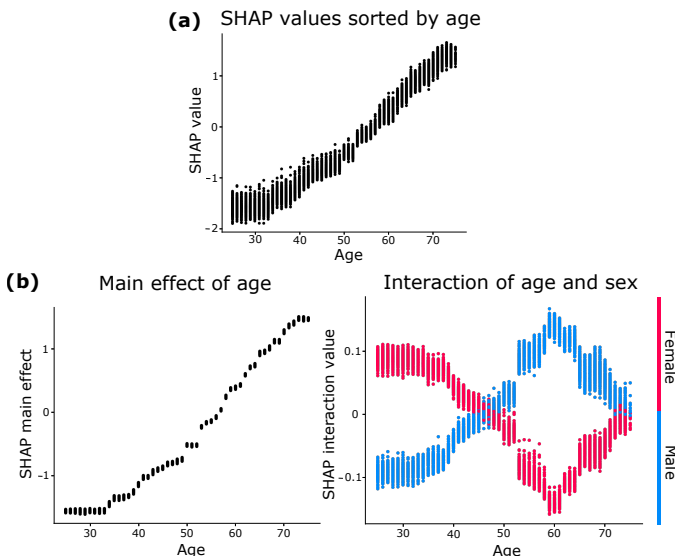


Fig. 17. Identifying mortality risk factors from explanations. (a) First-order SHAP values show that high age correlates with high mortality risk. The impact of other interacting factors on the age contribution to risk remain hidden. (b) The SHAP interaction matrix can be decomposed into the main effects (diagonal elements) and the interaction effects (off-diagonal elements). Here the interaction effect shows that the risk factor “age” is largely modulated by the sex of the participant. (Figure is adapted from [115].)

Fig. 17 (b) shows the main effect of age (left plot) and the interaction effect between age and sex (right plot) for the risk of mortality. These values constitute the diagonal and off-diagonal elements of the SHAP interaction matrix, respectively. One can see that age is on its own a very strong risk factor, however, depending on the sex of the person this risk factor is modulated differently over different age groups. While at a young age the mortality risk is slightly higher for women, it reverses and is much higher for men at age 60, and finally reaches the same level for the high age group. This example therefore shows that by only considering first-level explanations, we may miss important details and arrive at incomplete conclusions. Therefore, we recommend whenever possible to look at these interactions.

Look at interactions to get deeper insights

We note that a general limitation of higher-order explanations is that they need more data to be extracted with statistical significance. This is not a problem when considering models taking few input features and trained on large datasets, but this can become more challenging for high-dimensional data, and in that case, only first-order effects can be captured robustly. Furthermore, even with interaction effects, performing causal inference, e.g. deriving mortality risk factors from epidemiological data, remains difficult due to unobserved effects such as collider or measurement bias [69].

VIII. SUCCESSFUL USES OF EXPLANATION TECHNIQUES

Interpretation methods can be applied for a variety of purposes. Some works have aimed to understand the model’s prediction strategies, e.g. in order to validate the model [104]. Others visualize the learned representations and try to make the model itself more interpretable [75]. Finally, other works have sought to use explanations to learn about the data, e.g. by visualizing interesting input-prediction patterns extracted by a deep neural network model in scientific applications [186]. Technically, explanation methods have been applied to a broad range of models ranging from simple bag-of-words-type classifiers or logistic regression [13], [28] to feed-forward or recurrent deep neural networks [13], [171], [11], [9], and more recently also to unsupervised learning models [85], [86]. At the same time these methods were able to handle different types of data, including images [13], speech [20], text [10], [38], and structured data such as molecules [165], [162] or genetic sequences [191].

Some of the first successes in interpreting deep neural networks have occurred in the context of image classification, where deep convolutional networks have also demonstrated very high predictive performance [99], [65]. Explanation methods have for the first time allowed to open these “black boxes” and obtain insights into what the models have actually learned and how they arrive at their predictions. For instance, the works [174], [134]—also known in this context as “deep dreams”—highlighted surprising effects when analyzing the inner behavior of deep image classification models by synthesizing meaningful preferred stimuli. They report that the preferred stimuli for the class ‘dumbbell’ would indeed contain a visual rendering of a dumbbell, but the latter would systematically come with an arm attached to it [131], demonstrating that the output neurons do not only fire for the object of interest but also for correlated features.

Another surprising finding was reported in [101]. Here, interpretability—more precisely the ability to determine which pixels are being used for prediction—helped to reveal that the best performing ML model in a prestigious international competition, namely the PASCAL visual object classification (VOC) challenge, was actually relying partly on artefacts. The high performance of the model on the class “horse” could indeed be attributed to detecting a copyright tag present in the bottom left corner of many horse images of the dataset⁶, rather than detecting the actual horse in the image. Other

⁵In [115] the authors also propose a polynomial time algorithm for computing the Shapley value explanations on tree-based models.

⁶The presence of these artifacts in the benchmark dataset had gone unnoticed for almost a decade.

effects of similar type have been reported for other classes and datasets in many other works, e.g. in [147] models were shown to distinguish between the class “Husky” and “Wolf” solely based on the presence or absence of snow in the background.

These discoveries have been made rather accidentally by researchers carefully analysing suspicious explanations. It is clear that such laborious manual inspection of heatmaps does not scale to big datasets with *millions* of examples. Therefore, systematic approaches to the interpretation of ML models have recently gained increased attention.

A. From Explanations to Understanding Large ML Models

This section describes two examples of a systematic analysis of a large number of heatmaps. In the first case, the goal of the analysis is to systematically find data artefacts picked up by the model (e.g. copyright tags in horse images), whereas the second analysis aims to carefully investigate the learning process of a deep model, in particular the emergence of novel prediction strategies during training.

The process of systematically extracting data artefacts was automated by a method called Spectral Relevance Analysis (SpRay) [104], where after computing LRP-type explanations on a whole dataset (cf. Section VI-D), a cluster-based analysis was applied on the collection of produced explanations to extract *prototypical* decision strategies. The SpRay analysis would for example reveal for some shallow Fisher Vector model trained on Pascal VOC 2007 dataset that images of the ‘horse’ would be predicted as such using a finite number of prototypical decision strategies ranging from detecting the horse itself to detecting weakly related features such as horse racing poles, or clear artefacts such as copyright tags [101]. The analysis was later on applied to the decisions of a state-of-the-art VGG-16 deep neural network classifier trained on ImageNet, and here again, interesting insight about the decision structure could be identified [7]. Certain predictions, e.g. for the class ‘garbage truck’, could be found by SpRay to rely on some watermark in the bottom-left corner of the image (see Fig. 18). This watermark which is only present in specific images would thus be used by the model as a confounding factor (or artefact) to artificially improve prediction accuracy on this benchmark⁷.

Such strategy employed by the ML classifier can be referred to as ‘Clever Hans’ [104]. For machine learning models having implemented a Clever Hans strategy, an overconfident assessment of the true model accuracy would be produced by solely relying on the accuracy metric without an inspection of the model’s decision structure. The model would have likely performed erratically once it is applied in a real-world setting, where, e.g. the copyright tag is decoupled from the concept of a horse or garbage truck respectively. Here, the ability to explain the decision-making of the model and to automatically analyze these explanations on a very large dataset, was therefore a key ingredient to more robustly assess the model’s strength and weakness and potentially improving it.

⁷Or in the case of [7] *deteriorate* model performance, as the identified confounding feature is exclusive to the training data.

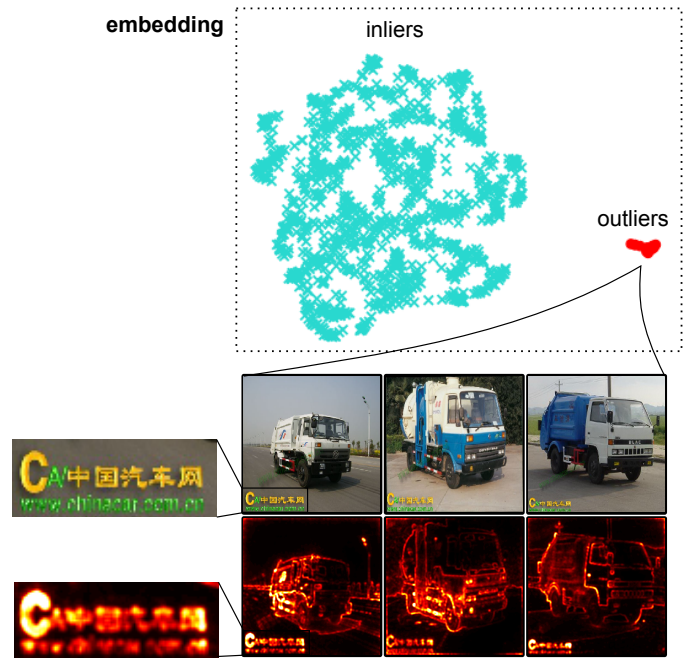


Fig. 18. SpRay analysis of the predictions of a pretrained VGG-16 model on images of the class ‘garbage truck’. *Top*: Low-dimensional embedding of the explained decisions for the class ‘garbage truck’. Points highlighted in red are outliers. *Bottom*: Images and corresponding decisions for some of the points highlighted in red.

Another example of a systematic interpretation of ML models can be found in the context of reinforcement learning, in particular board and video games. Here large amounts of data can be easily generated (simulated games) and used to carefully analyze the strategies of a ML model and how these strategies emerge during training. On games such as the arcade game Atari Breakout, the computer player would progressively learn strategies commonly employed by human players such as ‘tunnel-digging’ [124], [200]. The work of [104] analyzes the emergence of this advanced ‘tunnel-digging’ technique using explanations. First, LRP-type pixel-wise explanations of the player’s decision were produced at various time steps and training stages. The produced collection of explanations were then pooled (cf. Section VI-D1) on bounding boxes representing some key visual elements of the game, specifically, the ball, the paddle, and the tunnel. Pooled quantities could then be easily and quantitatively monitored over the different stages of training. The analysis is shown in Fig. 19.

We observe that the neural network model would first learn to play conventionally by keeping track of the ball and the paddle, and only at a later stage of the training process would learn to focus on the tunnel area, allowing the ball to go past the wall and bounce repeatedly in the top area of the screen. This analysis highlights in a way that is easy interpretable to the human the multi-stage nature of learning, in particular, how the learning machine progressively develops increasingly sophisticated game playing strategies. Overall, this summarized information on the decision structure of the model and on the evolution of the learning process could prove to be crucial in learning improved models on purposely consolidated datasets. They could also prove useful for characterizing the




Fig. 19. Analysis of the learning process of a deep model playing Atari Breakout. The curves show the development of the relative relevance of different game objects (ball, paddle, tunnel) averaged over six runs.

different stages of learning and developing more efficient training procedures.

B. From Explanations to Novel Scientific Insights

In the last subsection we demonstrated the use of explanation techniques for systematically analysing models and verifying that they have learned valid and meaningful prediction strategies. Once verified to not be Clever Hans predictors, non-linear models offer a lot of potential for the sciences to detect new interesting patterns in the data, which may lead to an improved understanding of the underlying natural structures and processes — the primary goal of scientists. So far this was not possible, because non-linear models were actually considered to be “black boxes”, i.e. scientists had to resort to the use of linear models (see e.g. [63], [117]), even if this came at the expense of predictivity. Only recently, the technologies have become available to extract scientific insights from complex nonlinear models [150], [166], [121]. In the following we will present a selection of problems where ML explanation techniques bring the full potential of non-linear methods to scientific disciplines.

Let us start with the discussion of scientific problems, which concern images. These problems could directly benefit from the advances made in general image recognition in the last years.

An important application area is in medicine [74]. Fig. 20 (a) shows one such application: the task of predicting tissue type from histopathology imagery. The work of [22] proposes a bag-of-words model for the prediction task with invariances to rotation, shift and scale of the input data. For the verification of the prediction results, the LRP technique is applied to this model so that heatmaps are produced, offering per-pixel scores which indicate the presence of tumorous structures. Furthermore, LRP heatmaps computed for different target cell types can be combined for obtaining computationally predicted fluorescence images. The explanations are histopathologically meaningful and may potentially give interesting information about which tissue components are most indicative of cancer. Further analyses such as the identification, localization and

counting of cells, i.e. lymphocytes, can be performed on these explanations (see [93]). Recently deeper models have also been used for predicting and explaining tissue type [61]. In addition to visual explanations, [204] also generate a free-text pathology report to clarify the decision of the classifier. Further applications of DNN explanation techniques for medical images include the detection of lesions in diabetic retinopathy data [145], pixel-wise analysis of microscopy images from global image annotations [98], the validation of predictions in dermatology [199], and analysis of x-ray images [48]. The latter work produces in addition to the image-level explanations, descriptive sentences to further improve the interpretability of the ML decision for a doctor.

Methods for explaining DNNs have also been applied outside of medical imaging, for example, in earth sciences. In [109], [43], the authors predict various meteorological values such as convection, or MRMS radar, from satellite imagery, and the latter prediction are then attributed to the individual pixels. An example is shown in Fig. 20 (b), where the MRMS output is supported by two visual pattern, the edge of the cloud, and a brighter spot representing a denser region of the cloud. Fig. 20 (c) shows a machine learning based analysis of galaxy morphologies [207]. The latter work aims to classify galaxy morphologies into five classes (completely round smooth, in-between smooth, cigar-shaped smooth, edge-on and spiral) using a convolutional neural networks, and the convolution filters as well as activation patterns are analyzed to understand what visual features are associated to these different classes (see Fig. 20 (f)). Lastly, the authors of [50] also rely on training and explaining DNN image classifiers to get novel insights, this time in the area of plant stress classification.

Explanation methods for deep neural networks have also demonstrated their potential beyond the image domain, e.g. on scientific problems concerning time series data. For instance, the work of [180] presents one of the first uses of DNNs explanation techniques in cognitive neurosciences, specifically in brain computer interfacing [40] where linear methods are still the most widely used filtering methods [24], [63]. The results in [180] show that deep models achieve similar decoding performances⁸ and learn neurophysiologically plausible patterns (see Fig. 20 (b)), namely focus on the contralateral sensorimotor cortex – an area where the event-related desynchronization occurs during motor imagery. However, in contrast to the patterns computed with conventional approaches [24], [63], which only allow to visualize the aggregated information (average activity) per class, the explanations computed with LRP are available for every single input of the deep learning classifier, i.e. for every time point of individual trials (see Fig. 20 (d)). This increased resolution (knowing which sources are relevant at each time point, instead of just having the average patterns) sheds further light onto cognitive processes in the brain.

Another application of DNN explanation techniques in cognitive neuroscience is presented in [186], where a deep model

⁸Deep models usually require larger amounts of training data to have an advantage over linear techniques.

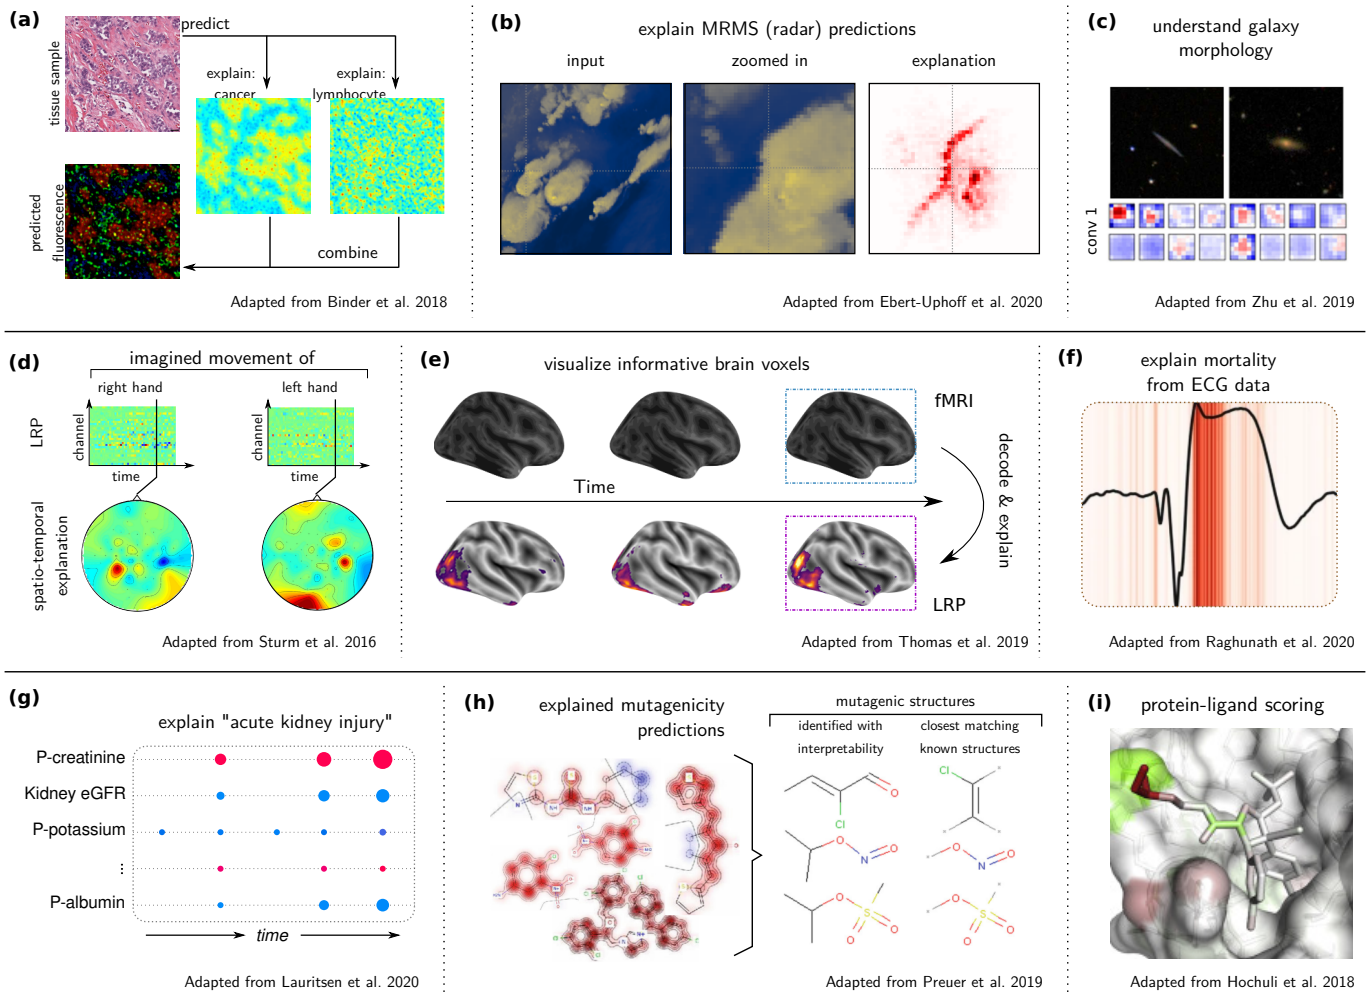


Fig. 20. Different applications of explanation techniques in the sciences. **(a)** LRP heatmaps merged into a computationally predicted fluorescence image. Here, red identifies cancer, green shows lymphocytes and blue is stroma. Adapted from [22]. **(b)** Prediction of MRMS radar signal from satellite image and pixel-wise explanation. Adapted from [43]. **(c)** Learned filter weights from the first convolutional layer of a deep neural network trained for galaxy morphology classification. Adapted from [207]. **(d)** A whole-brain fMRI volume is decoded using a DNN. The decoding decision is explained voxel-wise to localize brain areas corresponding to the predicted cognitive state. Adapted from [186]. **(e)** Example of LRP relevance maps for a single EEG trial of an imagined movement (each class). The matrices indicate the relevance of each time point (abscissa) and EEG channel (ordinate). Below the matrix the relevance information for two single time points is plotted as a scalp topography. Adapted from [180]. **(f)** Prediction of mortality from ECG time series data. Red areas indicate the time steps that most strongly explain the prediction. Adapted from [146]. **(g)** Some medical condition predicted from electronic health records (EHR) time series, and explained in terms of input features. Blue/red indicate low/high value and circle size indicates feature relevance. Adapted from [106]. **(h)** Graph convolutional neural prediction of the molecule's mutagenicity, and attribution of the on individual atoms. Interpretability feedback reveals that the model has correctly identified molecular substructures known to interact with (human) DNA. Adapted from [144]. **(i)** The predicted atom score describing protein-ligand interaction is explained with CLRP (green corresponding to a more favorable score). Adapted from [71].

is learned to predict whole-brain fMRI data. The method, termed DeepLight, outperforms well-established local or linear decoding methods such as the generalized linear model and searchlight (see [186]). An adaptation of LRP maintains interpretability and substantiates that the model's predictions are based on physiologically appropriate brain areas for the classified cognitive states. The approach is depicted in Fig. 20 (e) which visualizes the exemplar voxels that are used by the deep model to accurately decode the state from the fMRI signal. These voxels of high relevance have been shown to correspond very well to the active areas described in the fMRI literature (see [186]). Note that also here the deep model not only gives an advantage in terms of performance (i.e. better decoding accuracy) compared to the local or linear baseline

methods, but its explanations are provided for every single input, i.e. for every fMRI volume over time. This increased resolution allows to study the spatio-temporal dynamics of the fMRI signal and its impact on decoding, something which is not possible with classical decoding methods⁹.

Many other studies use explanation methods to analyze time series signals in the sciences. For example, [146] builds a model of mortality from ECG data, and applies DNN explanation techniques to identify for individual subjects, patterns in the ECG time series that support the prediction (cf. Fig. 20 (f) for one such explanation.) In another work

⁹In classical fMRI analyses, p-values indicate the relevance of brain voxels. However, these p-values are usually obtained on a subject or group-level, not for single trials or single time points.

[106], the input time series is formed by electronic health records (EHR) sampled at a much lower rate, from which a DNN model can decide among a set of medical conditions. The medical condition can then be attributed to the different medical measurements. An example of explanation is given in Fig. 20 (g). The produced explanation provides valuable additional information to decide in favor or against a particular treatment. Reference [77] introduces explanation techniques to the domain of human gait recognition and show that non-linear learning models are not only the better predictors but that they can at the same time learn physiologically meaningful features for subject prediction which align with expected features used by linear models. Another work [97] applies Long Short-Term Memory (LSTM) networks to the field of hydrology to predict the river discharge from meteorological observations. The authors apply the integrated gradients technique to analyze the internals of the network and obtain insights which are consistent with our understanding of the hydrological system.

Structured data such as molecules or gene sequences are another very important domain for scientific investigations. Therefore, non-linear ML coupled with explanation techniques have also attracted attention in scientific communities working with this type of data. One successful example of the use of DNN explanation methods in this domain has been reported in [144]. The authors train a deep model to predict molecular properties and bioactivities and apply the Integrated Gradients method to analyze what the model has learned and extract interesting insights (see Fig. 20 (h)). For instance, they find that single neurons play the role of pharmacophore detectors and demonstrate that the model uses pharmacophore-like features to reach its conclusions, which are consistent with pharmacologist literature. Another work [71] (see Fig. 20 (i)) applies an extended version of LRP called CLRP to visualize how CNNs interpret individual protein-ligand complexes in molecular modeling. Also here the trained model learns meaningful features and has the ability to provide new insights into the mechanisms underlying protein-ligand interactions. Yet another work [196] applies LSTM predictors together with LRP for transparent therapy prediction on patients suffering from metastatic breast cancer. Clinical experts verify that the features used for prediction, as revealed via LRP, largely agree with established clinical guidelines and knowledge. The work by [87] uses a gradient-based explanation technique to understand the activity prediction across chromosomes, whereas [31] uses the LRP explanation technique for understanding automated decisions on behavioral biometrics. Recently, also the physics community started to use machine learning explanations for the task of energy prediction. The work of [165], [166] showed that accurate predictions are possible and obtained also physical meaningful insights from the model. Other works [114] showed that explanations in gene analysis lead to interpretable patterns consistent with literature knowledge.

IX. CHALLENGES AND OUTLOOK

While recent years have seen astonishing conceptual and technical progress in XAI, it is important to carefully discuss

the current limits and the challenges that will need to be addressed by researchers to further establish the field and increase the usefulness of XAI systems.

Foundational theoretical work in XAI has so far been limited. As discussed above in Section V, Shapley values [169], [116], and (deep) Taylor decomposition [127] have been proposed as principled frameworks for formalizing the task of explanation. Other frameworks such as rate distortion theory [118] have also been proposed. Numerous theoretical questions however remain: For example, it remains unclear how to weigh the model and the data distribution into the explanation, in particular, whether an explanation should be based on *any* features the model locally reacts to, or only those that are expressed locally. Related to this question is that of causality, i.e. assuming a causal link between two input variables, it has not been answered yet whether the two variables, or only the source variable, must constitute the explanation. A deeper formalization and theoretical understanding of XAI will be instrumental for shedding light into these important questions.

Another central question in XAI is that of optimality of an explanation. So far, there is no well-agreed understanding of what should be an optimal explanation. Also, ground-truth explanations cannot be collected by humans as this would presuppose they are able to make sense of the complex ML model they would like to explain in the first place. Methods such as ‘pixel-flipping’ [158] assess explanation quality indirectly by testing how flipping relevant pixels affects the output score. The ‘axiomatic approach’ [183], [125] does not have this indirect step, however, axioms are usually too generic to evaluate an explanation comprehensively. Approaches relying on ground truth information (e.g. [12]) derived from a synthetic dataset offer a direct and objective way to evaluate and compare explanations, however, it needs to be demonstrated that the synthetic problem is representative of the typically more complex real-world problem, where the ground truth explanations are usually not available. Self-explainable models (e.g. [154]) provide another potential avenue towards optimal explanations, by forcing the model to be built in a way that explanations can be unambiguously extracted, thereby bypassing some of the technical challenges of post-hoc explanations. Such self-explainable approach comes however at the possible expense of prediction accuracy or runtime efficiency. Hence, the question of producing optimal explanations is still largely an open question. Lastly, accuracy is only one factor in a more general assessment of the overall practical value of an explanation, which further includes human factors such as understandability, manageability, and overall utility of the XAI system [184], [147], [133]. Application-driven evaluations account for these additional factors, however, they are also hard to implement in practice [41].

Further challenges arise when applying XAI on problems where different actors (e.g. the explainer and the explainee) have conflicting interests. Recent work has shown that an ‘adversary’ can modify the ML model in an imperceptible fashion so that the prediction behavior remains intact but the explanation of those predictions changes drastically [68]. Relatedly, even when the model remains unchanged, inputs could be perturbed imperceptibly to produce arbitrary explana-

tions [39]. An important challenge will be to provide provable guarantees on the robustness of explanations to various types of external distortion.

Interpretability may also find itself at odds with the constant quest for higher predicting accuracy. The model development and the explanation of an already trained model are often treated as two independent processes, i.e. the model developers aim to build the best possible model and the post-hoc XAI community provides the tools to interpret it. Because highly predictive models are becoming increasingly complex both in terms of number of parameters and structure of the model, XAI software must keep up with this increasing complexity [3], and at the same time, the human must also deal with explanations of increasingly subtler predictions [14]. Same challenges also occur for self-explainable models (cf. Section III-E), where the incorporated interpretability structures, would have to be continuously refined and extended to cope with the increased complexity of data and tasks, and the intrinsic limits of the human receiving the explanations.

When designing new XAI-driven applications, adopting a holistic view that sets the right tradeoffs and delivers the optimal amount of information and range of action to the multiple and potentially conflicting actors, will constitute an important practical challenge.

Another question of utmost importance, especially, in safety critical domains, is whether we can fully trust the model after having explained some predictions. Here, we need to distinguish between model interpretation and model certification: While it is helpful to explain models for available input data, e.g. to potentially detect some erroneous decision strategies, certification would require to verify the model for *all possible* inputs, not only those included in the data. Furthermore, it must be remembered that explanations returned to the user are summaries of a potentially complex decision process, i.e. there may be different decision strategies, the wrong ones and the correct ones, mapping to the same explanation. Lastly, explanations are subject to their own biases and approximations, and they can be manipulated by an adversary to loose their informative content. Therefore, in order to ultimately establish a truly safe and trustworthy model, further steps are needed, potentially including the use of formal verification methods [21], [84].

Moreover, it may be worthwhile to explore new forms of explanations that are optimally suited to their user. Such explanations could for example leverage the user’s prior knowledge or personal preferences. Novel approaches from knowledge engineering, cognitive sciences, and human-computer interfaces, will need to contribute. Also, while heatmaps provide a first intuition to users, they may not take advantage of the complex abstract reasoning capabilities of humans and can be very difficult to interpret for certain data modalities (e.g. time series data). As we have discussed in the second worked-through example, explaining using a different modality (or by projecting the data on a more interpretable basis) can lead to an explanation that is more useful to the user. An important challenge will be to develop systematic ways of transferring explanations from the original input domain to a more interpretable target domain.

Finally, we see as a key future challenge the design of XAI techniques that can automatically extract meaningful collective variables and explain in their terms. Collective variables are central in many area of physics. In solid state physics, they have led to groundbreaking advances, defining quasiparticles such as phonons, plasmons, polarons, magnons, excitons [92], etc. Ideally, collective variables in this sense would be inferred from a learning model by e.g. automatically binding explanation variables in meaningful abstract ways. In the neurosciences, von der Malsburg has coined the concept of ‘binding’ for neural strategies that allow sets of variables (neurons) to synchronize collectively by learning [192]. While steps have already been taken in XAI to identify interacting variables [188], [34], [42], [162], [115], it will be necessary for a broader usage to generalize the approach, e.g. to extract ‘mathematical formulas’ that first build the needed collective variables, and use them to explain, concisely but deeply, the ML predictions.

X. CONCLUSION

Complex nonlinear ML models such as neural networks or kernel machines have become game changers in the sciences and industry. Fast progress in the field of explainable AI, has made virtually any of these complex models, supervised or unsupervised, interpretable to the user. Consequently, we no longer need to give up predictivity in favor of interpretability, and we can take full advantage of strong nonlinear machine learning in practical applications.

In this review we have made the attempt to provide a systematic path to bring XAI to the attention of an interested readership. This included an introduction to the technical foundations of XAI, a presentation of practical algorithms such as Occlusion, Integrated Gradients and LRP, concrete examples illustrating how to use explanation techniques in practice, and a discussion of successful applications. We would like to stress that the techniques introduced in this paper can be readily and broadly applied to the workhorses of supervised and unsupervised learning, e.g. clustering, anomaly detection, kernel machines, deep networks, as well as state-of-the-art pretrained convolutional networks and LSTMs.

XAI techniques not only shed light into the inner workings of non-linear learning machines, explaining why they arrive at their successful predictions; they also help to discover biases and quality issues in large data corpora with millions of examples [7]. This is an increasingly relevant direction since modern machine learning relies more and more on reference datasets and reference pretrained models. Furthermore, initial steps have been taken to use XAI beyond validation to arrive at better and more predictive models e.g. [152], [6], [8], [7].

We would like to stress the importance of XAI, notably in safety critical operations such as medical assistance or diagnosis, where the highest level of transparency is required in order to avoid fatal outcomes.

Finally as a versatile tool in the sciences, XAI has been allowing to gain novel insights (e.g. [165], [22], [71], [186], [45], [159], [140]) ultimately contributing to further our scientific knowledge.

While XAI has seen an almost exponential rise in interest (and progress) with communities forming and many workshops emerging, there is a wealth of open problems and challenges with ample opportunities to contribute (see Section IX). Concluding, we firmly believe that XAI will in the future become an indispensable practical ingredient to obtain improved, transparent, safe, fair and unbiased learning models.

ACKNOWLEDGMENT

This work was supported by the Institute of Information & Communications Technology Planning & Evaluation (IITP) grants funded by the Korea Government (No. 2017-0-00451, Development of BCI based Brain and Cognitive Computing Technology for Recognizing User’s Intentions using Deep Learning and No. 2019-0-00079, Artificial Intelligence Graduate School Program, Korea University), by the German Ministry for Education and Research (BMBF) under Grants 01IS14013A-E, 01GQ1115, 01GQ0850, 01IS18025A and 01IS18037A; and by the German Research Foundation (DFG) under Grant Math+, EXC 2046/1, Project ID 390685689. Correspondence to WS, GM, KRM.

APPENDIX A

IMPLEMENTING SMOOTH INTEGRATED GRADIENTS

In this appendix, we give the algorithm combining SmoothGrad [176] and Integrated Gradients [183], which we use in Section IV in our comparison of explanation methods. Its implementation is shown in Algorithm 5.

Algorithm 5 Integrated Gradients with Smoothing

```

R = 0
for  $s = 1 \dots S$  do
   $\tilde{x} \sim \mathcal{N}(0, \sigma I)$ 
  for  $t = 1 \dots T$  do
     $\mathbf{R} = \mathbf{R} + (\mathbf{x} - \tilde{x}) \odot \nabla f(\tilde{x} + \frac{t-0.5}{T} \cdot (\mathbf{x} - \tilde{x}))$ 
  end for
end for
return  $\frac{1}{TS} \cdot \mathbf{R}$ 

```

The procedure consists of a simple nested loop of S smoothing and T integration steps, where each integration starts at some random location near the origin. Here, we note that these locations are not strict root points. However, in the context of image data, random noise does not change significantly evidence in favor or against a particular class. Thus, the explanation remains approximately complete.

APPENDIX B

IMPLEMENTING LAYER-WISE RELEVANCE PROPAGATION

In this appendix, we outline two possible implementations of LRP [13], [126]. A first one that is intuitive and based on looping forward and backward over the multiple layers of the neural network. This procedure can be applied to simple sequential structures such as VGG-16 [175]. The second approach we present is based on ‘forward hooks’ and serves to extend the LRP method to more complex architectures such as ResNet [65].

A. Standard Sequential Implementation

The standard implementation is based on the forward-backward procedure outlined in Algorithm 2. We focus here on the `relprop` function of this procedure, which is called at each layer to propagate relevance to the layer below. We give an implementation for the LRP-0/ ϵ/γ rules [13], [126] and one for the z^B -rule [127]. The first three rules can be seen as special cases of the more general rule

$$R_j = \sum_k \frac{a_j \rho(w_{jk})}{\epsilon + \sum_{0,j} a_j \rho(w_{jk})} R_k$$

where $\rho(w_{jk}) = w_{jk} + \gamma w_{jk}^+$. This propagation rule can be computed in four steps.

Algorithm 6 LRP-0/ ϵ/γ

```

 $\mathbf{z} = \epsilon + f_l^\rho(\mathbf{a}^{(l-1)})$  (Step 1)
 $\mathbf{s} = \mathbf{R}^{(l)} \odot \mathbf{z}$  (Step 2)
 $\mathbf{c} = \nabla \langle \mathbf{z}, [\mathbf{s}]_{\text{cst.}} \rangle$  (Step 3)
 $\mathbf{R}^{(l-1)} = \mathbf{a} \odot \mathbf{c}$  (Step 4)
return  $\mathbf{R}^{(l-1)}$ 

```

The first step applies f_l^ρ , a forward evaluation of a copy of the layer whose parameters have gone through some function ρ , and also adds a small positive term ϵ . The third step is conveniently expressed as a gradient of some dot product $\langle \mathbf{z}, [\mathbf{s}]_{\text{cst.}} \rangle$ w.r.t. the input activations. The notation $[\cdot]_{\text{cst.}}$ indicates that the term has been detached from the gradient computation and is therefore treated as a constant. In PyTorch, for example, this can be achieved by calling `() .data`. The `relprop` function implemented by Algorithm 6 is applicable for most linear and convolution layers of a deep rectifier network. For the pixel-layer, we use instead the z^B -rule [127], [126]:

$$R_i = \sum_j \frac{x_i w_{ij} - l_i w_{ij}^+ - h_i w_{ij}^-}{\sum_i x_i w_{ij} - l_i w_{ij}^+ - h_i w_{ij}^-} R_j$$

where l_i and h_i are the lowest/highest possible pixel values of x_i , and where $w_{ij}^- = \min(0, w_{ij})$. The corresponding implementation is shown in Algorithm 7 and again consists of four steps:

Algorithm 7 z^B -rule

```

 $\mathbf{z} = f_1^+(\mathbf{x}) - f_1^+(\mathbf{l}) - f_1^-(\mathbf{h})$  (Step 1)
 $\mathbf{s} = \mathbf{R}^{(1)} \odot \mathbf{z}$  (Step 2)
 $\mathbf{c} = \nabla_{\mathbf{x}, \mathbf{l}, \mathbf{h}} \langle \mathbf{z}, [\mathbf{s}]_{\text{cst.}} \rangle$  (Step 3)
 $\mathbf{R}^{(0)} = \mathbf{x} \odot \mathbf{c}_1 + \mathbf{l} \odot \mathbf{c}_2 + \mathbf{h} \odot \mathbf{c}_3$  (Step 4)
return  $\mathbf{R}^{(0)}$ 

```

The functions f_1^+ and f_1^- are forward passes on copies of the first layer whose parameters have been processed by the functions $\max(0, \cdot)$ and $\min(0, \cdot)$ respectively.

B. Forward-Hook Implementation

When the architecture has non-sequential components (e.g. ResNet [65]), it is more convenient to reuse the graph traversing procedures readily implemented by the model’s

existing forward pass and the automatically generated gradient propagation pass. To achieve this, we can implement ‘forward hooks’ at each linear and convolution layers. In this case, we leverage the ‘smooth gradient’ view of LRP (cf. Eq. (3)) and modify the implementation of the forward pass in a way that it keeps the forward pass functionally equivalent but modifies the local gradient computation. This is achieved by strategically detaching terms from the gradient in a way that calling the gradient becomes equivalent to computing Eq. (3) at each layer. Once the forward functions have been redefined at each layer, the explanation can be computed globally by calling the gradient of the whole function as shown in Algorithm 8. (Note that unlike the original function $f(\mathbf{x})$ the new function that includes the hooks receives three arguments as input: the data point \mathbf{x} , and the bounds \mathbf{l} and \mathbf{h} used by the first layer.)

Algorithm 8 LRP implementation based on forward hooks

Forward hook for intermediate layers (LRP-0/ ϵ/γ)

```

 $z = \epsilon + f_i^p(\mathbf{a}^{(l-1)})$ 
return  $z \odot [f_i(\mathbf{a}^{(l-1)}) \oslash z]_{\text{cst.}}$ 

```

Forward hook for the first layer (z^B -rule)

```

 $z = f_1(\mathbf{x}) - f_1^+(\mathbf{l}) - f_1^-(\mathbf{h})$ 
return  $z \odot [f_1(\mathbf{x}) \oslash z]_{\text{cst.}}$ 

```

Global LRP computation

```

 $y = f(\mathbf{x}, \mathbf{l}, \mathbf{h})$ 
 $\mathbf{c}_1, \mathbf{c}_2, \mathbf{c}_3 = \hat{\nabla} y$ 
 $\mathbf{R} = \mathbf{x} \odot \mathbf{c}_1 + \mathbf{l} \odot \mathbf{c}_2 + \mathbf{h} \odot \mathbf{c}_3$ 
return  $\mathbf{R}$ 

```

The forward-hook implementation produces exactly the same output as the original function $f(\mathbf{x})$, but its ‘gradient’, which we denote by $\hat{\nabla}$ is no longer the same due to the detached terms. As a result, calling the gradient of this function, and recombining it with the input yields the same desired LRP explanation as one would get with the standard LRP implementation, but has now gained applicability to a broader set of neural network architectures.

APPENDIX C

EXPLANATION SOFTWARE

The attention to interpretability in machine learning has grown frantically throughout the past decade alongside research on, and the development of computationally efficient deep learning frameworks. This attention in turn caused a strong demand for accessible and efficient software solutions for out-of-the-box applicability of XAI. In this section we briefly highlight a collection of software toolboxes released in recent years, providing convenient access to a plethora of methods of XAI and supporting various computational backends. A summarizing overview over the presented software solutions is given in Table II, alongside a glossary of methods with respective abbreviations used throughout our review in Table III.

One of the earlier and comprehensive XAI software packages is the LRP Toolbox [102], providing presently up to date implementations of LRP for the — until very recently — popular Caffe deep learning framework [83], as well as Matlab and Python via custom neural network interfaces. While support for Caffe is restricted to the C++ programming language and thus to CPU hardware, it provides functionality implementing DCN, GB, DTD, and SA and can be built and used as a stand-alone executable binary for predictors based on the Caffe neural network format. The sub-packages available for Matlab and Python provide out-of-the-box support for LRP and SA, while being easily extensible via custom neural network modules written with clarity and the methods’ intelligibility in mind. The `cupy` [137] backend constitutes an alternative to the CPU-bound `numpy` [138] package, providing optional support for modern GPU hardware from NVIDIA.

Both the DeepExplain [4] and iNNvestigate [3] toolboxes built on top of the popular Keras [30] package for Python with TensorFlow backend for explaining Deep Neural Network models, and thus provide support for both CPU and GPU hardware and convenient access for users of Keras models. While the more recent iNNvestigate Toolbox implements a superset of the modified backpropagation methods available in DeepExplain, the latter also offers functionality for perturbation-based attribution methods, i.e. the Occlusion method [201] and Shapley Value Resampling [29]. For explaining a model’s prediction DeepExplain allows for an ad-hoc selection of the explanation method via pythonic context managers. The iNNvestigate package on the other hand operates by attaching and automatically configuring (several) modified backward graphs called “analyzers” to a model of interest — one per XAI method to compute attributions with.

A present trend in the machine learning community is a migration to the PyTorch framework with its eager execution paradigm, away from other backends. Both the TorchRay [46] and Captum [95] packages for Python and PyTorch enable the use of interpretability methods for neural network models defined in context of PyTorch’s high level neural network description modules. Captum can be understood as a rich selection of XAI methods based on modified backprop and is part of the PyTorch project itself. While not as extensive as Captum, the TorchRay package offers a series benchmarks for XAI alongside its selection of (benchmarked) interpretability methods.

REFERENCES

- [1] J. Adebayo, J. Gilmer, I. J. Goodfellow, and B. Kim. Local explanation methods for deep neural networks lack sensitivity to parameter values. In *International Conference on Learning Representations*, 2018.
- [2] C. Agarwal and A. Nguyen. Explaining image classifiers by removing input features using generative models. In *Proceedings of the Asian Conference on Computer Vision*, 2020.
- [3] M. Alber, S. Lapuschkin, P. Seegerer, M. Hägele, K. T. Schütt, G. Montavon, W. Samek, K.-R. Müller, S. Dähne, and P.-J. Kindermans. iNNvestigate neural networks! *Journal of Machine Learning Research*, 20:93:1–93:8, 2019.
- [4] M. Ancona, E. Ceolini, C. Öztireli, and M. Gross. Towards better understanding of gradient-based attribution methods for deep neural networks. In *International Conference of Learning Representations (ICLR)*, 2018.

TABLE II
INTERPRETABILITY SOFTWARE PACKAGES BY TIME OF RELEASE

Software Package	Release	Available from https://github.com/ ...	Compute Backend	GPU Support	Methods
LRP Toolbox [102]	2016	sebastian-lapuschkin/lrp_toolbox	Caffe numpy/cupy Matlab	✗ ✓ ✗	DCN, DTD, GB, LRP, SA LRP, SA LRP, SA
DeepExplain [4]	2017	marcoancona/DeepExplain	Keras+TensorFlow	✓	DLR, IG, LRP- ϵ SA, AS, OCC
iNNvestigate [3]	2019	albermax/innvestigate	Keras+TensorFlow	✓	DCN, DL, DTD, GB, IG, LRP, PA, PN, SA, SG Perturbation Analysis
TorchRay [46]	2019	facebookresearch/TorchRay	PyTorch	✓	DCN, EB, EP GC, GB, RISE, SA, various benchmarks
Captum [95]	2019 (beta)	pytorch/captum	PyTorch	✓	DCN, DLR, DLSHAP, GB, GC, GGC, GSHAP, IG, II, LC, NC, NGB, NIG, SA, SG, SG-SQ, TC, VG

TABLE III
GLOSSARY OF INTERPRETABILITY METHODS WITH ABBREVIATIONS REFERENCED THROUGHOUT OUR REVIEW

Method	Abbrev.	Method	Abbrev.
Anchors [148]	ANCH	Layer-wise Relevance Propagation (full) [13]	LRP
ApproShapley (Shapley Value Sampling) [29]	AS	LRP (composite strategy) [103], [126]	LRP-CMP
Class Activation Mapping [206]	CAM	LRP (specific variants) [13], [126]	LRP-*
Contextual Prediction Difference Analysis [56]	CPDA	Local Interpretable Model-agnostic Explanations [147]	LIME
DeconvNet [201]	DCN	Meaningful Perturbation [47]	MP
DeepLIFT [170]	DL	NeuronConductance [36]	NC
DeepLIFT (Rescale) [170]	DLR	NeuronGuidedBackprop [178]	NGB
DeepLIFT SHAP [116]	DLSHAP	NeuronIntegratedGradients [172]	NIG
Deep Taylor Decomposition [127]	DTD	Occlusion Analysis [201]	OCC
ExcitationBackprop [202]	EB	PatternAttribution [90]	PA
ExtremalPerturbation [46]	EP	PatternNet [90]	PN
GNNExplainer [198]	GNNEXP	Prediction Difference Analysis [208]	PDA
GNN-LRP [162]	GLRP	Randomized Input Sampling for Explanation [142]	RISE
GradCAM [167]	GC	Saliency Analysis / Gradient [14], [174]	SA
Gradient SHAP [116]	GSHAP	SHapley Additive exPlanations [116]	SHAP
Gradient \times Input [170]	GI	SHAP Interaction Index [115]	SHAPIDX
GuidedBackprop [178]	GB	SmoothGrad [176]	SG
Guided GradCam [167]	GGC	SmoothGrad ² [76]	SG-SQ
Integrated Gradients [183]	IG	Spectral Relevance Analysis [104]	SpRay
Internal Influence [110]	II	TreeExplainer [115]	TEXP
Kernel SHAP [116]	KSHAP	VarGrad [1]	VG
LayerConductance [172]	LC	Testing with Concept Activation Vectors [89]	TCAV
Local Rule-based Explanations [58]	LORE	TotalConductance [36]	TC

- [5] M. Ancona, C. Oztireli, and M. Gross. Explaining deep neural networks with a polynomial time algorithm for shapley value approximation. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97, pages 272–281, 2019.
- [6] C. J. Anders, G. Montavon, W. Samek, and K.-R. Müller. Understanding patch-based learning of video data by explaining predictions. In *Explainable AI: Interpreting, Explaining and Visualizing Deep Learning*, pages 297–309. Springer, 2019.
- [7] C. J. Anders, L. Weber, D. Neumann, W. Samek, K.-R. Müller, and S. Lapuschkin. Finding and removing Clever Hans: Using explanation methods to debug and improve deep models. *arXiv preprint arXiv:1912.11425*, 2019.
- [8] J. A. Arjona-Medina, M. Gillhofer, M. Widrich, T. Unterthiner, J. Brandstetter, and S. Hochreiter. Rudder: Return decomposition for delayed rewards. In *Advances in Neural Information Processing Systems*, pages 13544–13555, 2019.
- [9] L. Arras, J. A. Arjona-Medina, M. Widrich, G. Montavon, M. Gillhofer, K.-R. Müller, S. Hochreiter, and W. Samek. Explaining and interpreting lstms. In *Explainable AI: Interpreting, Explaining and Visualizing Deep Learning*, volume 11700 of *Lecture Notes in Computer Science*, pages 211–238. Springer, 2019.
- [10] L. Arras, F. Horn, G. Montavon, K.-R. Müller, and W. Samek. "What is relevant in a text document?": An interpretable machine learning approach. *PLoS ONE*, 12(8):e0181142, 2017.
- [11] L. Arras, G. Montavon, K.-R. Müller, and W. Samek. Explaining recurrent neural network predictions in sentiment analysis. In *Proceedings of the EMNLP'17 Workshop on Computational Approaches to Subjectivity, Sentiment & Social Media Analysis (WASSA)*, pages 159–168, 2017.
- [12] L. Arras, A. Osman, and W. Samek. Ground truth evaluation of neural network explanations with clevr-xai. *arXiv:2003.07258*, 2020.
- [13] S. Bach, A. Binder, G. Montavon, F. Klauschen, K.-R. Müller, and W. Samek. On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation. *PLoS ONE*, 10(7):e0130140, 2015.
- [14] D. Baehrens, T. Schroeter, S. Harmeling, M. Kawanabe, K. Hansen, and K.-R. Müller. How to explain individual classification decisions. *Journal of Machine Learning Research*, 11:1803–1831, 2010.
- [15] D. Bahdanau, K. Cho, and Y. Bengio. Neural machine translation by jointly learning to align and translate. In *3rd International Conference on Learning Representations*, 2015.
- [16] D. Balduzzi, M. Frean, L. Leary, J. P. Lewis, K. W. Ma, and B. McWilliams. The shattered gradients problem: If resnets are the answer, then what is the question? In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 342–350. PMLR, 2017.

- [17] D. Bau, J. Zhu, H. Strobelt, B. Zhou, J. B. Tenenbaum, W. T. Freeman, and A. Torralba. GAN dissection: Visualizing and understanding generative adversarial networks. In *7th International Conference on Learning Representations*, 2019.
- [18] D. Bau, J.-Y. Zhu, H. Strobelt, A. Lapedriza, B. Zhou, and A. Torralba. Understanding the role of individual units in a deep neural network. *Proceedings of the National Academy of Sciences*, 117(48):30071–30078, 2020.
- [19] S. Bazen and X. Joutard. The Taylor decomposition: A unified generalization of the Oaxaca method to nonlinear models. Working papers, HAL, 2013.
- [20] S. Becker, M. Ackermann, S. Lapuschkin, K.-R. Müller, and W. Samek. Interpreting and explaining deep neural networks for classification of audio signals. *arXiv preprint arXiv:1807.03418*, 2018.
- [21] F. Berkenkamp, M. Turchetta, A. Schoellig, and A. Krause. Safe model-based reinforcement learning with stability guarantees. In *Advances in Neural Information Processing Systems*, pages 908–918, 2017.
- [22] A. Binder, M. Bockmayr, M. Hägele, S. Wienert, D. Heim, K. Hellweg, A. Stenzinger, L. Parlow, J. Budczies, B. Goepfert, D. Treue, M. Kotani, M. Ishii, M. Dietel, A. Hocke, C. Denkert, K.-R. Müller, and F. Klauschen. Morphological and molecular breast cancer profiling through explainable machine learning. *Nature Machine Intelligence*, 2021.
- [23] C. M. Bishop. *Neural Networks for Pattern Recognition*. Oxford University Press, Inc., USA, 1996.
- [24] B. Blankertz, R. Tomioka, S. Lemm, M. Kawanabe, and K.-R. Müller. Optimizing spatial filters for robust eeg single-trial analysis. *IEEE Signal Processing Magazine*, 25(1):41–56, 2008.
- [25] W. Brendel and M. Bethge. Approximating cnns with bag-of-local-features models works surprisingly well on imagenet. In *7th International Conference on Learning Representations*, 2019.
- [26] N. Carlini and D. Wagner. Towards evaluating the robustness of neural networks. In *2017 IEEE Symposium on Security and Privacy (SP)*, pages 39–57. IEEE, 2017.
- [27] M. Caron, P. Bojanowski, A. Joulin, and M. Douze. Deep clustering for unsupervised learning of visual features. In *15th European Conference on Computer Vision*, volume 11218 of *Lecture Notes in Computer Science*, pages 139–156. Springer, 2018.
- [28] R. Caruana, Y. Lou, J. Gehrke, P. Koch, M. Sturm, and N. Elhadad. Intelligible models for healthcare: Predicting pneumonia risk and hospital 30-day readmission. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1721–1730, 2015.
- [29] J. Castro, D. Gómez, and J. Tejada. Polynomial calculation of the shapley value based on sampling. *Computers & Operations Research*, 36(5):1726–1730, 2009.
- [30] F. Chollet et al. Keras. <https://keras.io>, 2015.
- [31] P. Chong, Y. X. M. Tan, J. Guarnizo, Y. Elovici, and A. Binder. Mouse authentication without the temporal aspect – what does a 2d-cnn learn? In *2018 IEEE Security and Privacy Workshops (SPW)*, pages 15–21. IEEE, 2018.
- [32] C. S. Cox. *Plan and operation of the NHANES I Epidemiologic Followup Study, 1987*. Number 27. US Department of Health and Human Services, Public Health Service, Centers for Disease Control, National Center for Health Statistics, 1992.
- [33] D. R. Cox. Regression models and life-tables. *Journal of the Royal Statistical Society: Series B (Methodological)*, 34(2):187–202, 1972.
- [34] T. Cui, P. Marttinen, and S. Kaski. Learning global pairwise interactions with bayesian neural networks. In *24th European Conference on Artificial Intelligence*, volume 325 of *Frontiers in Artificial Intelligence and Applications*, pages 1087–1094. IOS Press, 2020.
- [35] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 248–255, 2009.
- [36] K. Dhamdhere, M. Sundararajan, and Q. Yan. How important is a neuron? *arXiv preprint arXiv:1805.12233*, 2018.
- [37] I. S. Dhillon, Y. Guan, and B. Kulis. Kernel k-means: spectral clustering and normalized cuts. In *Proceedings of the Tenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 551–556. ACM, 2004.
- [38] Y. Ding, Y. Liu, H. Luan, and M. Sun. Visualizing and understanding neural machine translation. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*, pages 1150–1159, 2017.
- [39] A. Dombrowski, M. Alber, C. J. Anders, M. Ackermann, K.-R. Müller, and P. Kessel. Explanations can be manipulated and geometry is to blame. In *Advances in Neural Information Processing Systems*, pages 13567–13578, 2019.
- [40] G. Dornhege, J. d. R. Millán, T. Hinterberger, D. McFarland, K.-R. Müller, et al. *Toward Brain-Computer Interfacing*, volume 63. MIT press Cambridge, MA, 2007.
- [41] F. Doshi-Velez and B. Kim. A roadmap for a rigorous science of interpretability. *arXiv preprint arXiv:1702.08608*, 2017.
- [42] O. Eberle, J. Büttner, F. Krätli, K.-R. Müller, M. Valleriani, and G. Montavon. Building and interpreting deep similarity models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 1–1, 2020.
- [43] I. Ebert-Uphoff and K. Hilburn. Evaluation, tuning, and interpretation of neural networks for working with images in meteorological applications. *Bulletin of the American Meteorological Society*, 101(12):E2149–E2170, 2020.
- [44] E. Eidinger, R. Enbar, and T. Hassner. Age and gender estimation of unfiltered faces. *IEEE Transactions on Information Forensics and Security*, 9(12):2170–2179, 2014.
- [45] H. J. Escalante, S. Escalera, I. Guyon, X. Baró, Y. Güçlütürk, U. Güçlü, and M. van Gerven. *Explainable and Interpretable Models in Computer Vision and Machine Learning*. Springer, 2018.
- [46] R. Fong, M. Patrick, and A. Vedaldi. Understanding deep networks via extremal perturbations and smooth masks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2950–2958, 2019.
- [47] R. C. Fong and A. Vedaldi. Interpretable explanations of black boxes by meaningful perturbation. In *IEEE International Conference on Computer Vision (ICCV)*, pages 3449–3457, 2017.
- [48] W. Gale, L. Oakden-Rayner, G. Carneiro, L. J. Palmer, and A. P. Bradley. Producing radiologist-quality reports for interpretable deep learning. In *16th IEEE International Symposium on Biomedical Imaging*, pages 1275–1279, 2019.
- [49] J. Gama, I. Zliobaite, A. Bifet, M. Pechenizkiy, and A. Bouchachia. A survey on concept drift adaptation. *ACM Computing Surveys*, 46(4):44:1–44:37, 2014.
- [50] S. Ghosal, D. Blystone, A. K. Singh, B. Ganapathysubramanian, A. Singh, and S. Sarkar. An explainable deep machine vision framework for plant stress phenotyping. *Proceedings of the National Academy of Sciences*, 115(18):4613–4618, 2018.
- [51] X. Glorot, A. Bordes, and Y. Bengio. Deep sparse rectifier neural networks. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, volume 15 of *JMLR Proceedings*, pages 315–323. JMLR.org, 2011.
- [52] I. Goodfellow, Y. Bengio, and A. Courville. *Deep learning*. MIT press, 2016.
- [53] I. J. Goodfellow, J. Shlens, and C. Szegedy. Explaining and harnessing adversarial examples. In *3rd International Conference on Learning Representations*, 2015.
- [54] B. Goodman and S. R. Flaxman. European union regulations on algorithmic decision-making and a “right to explanation”. *AI Magazine*, 38(3):50–57, 2017.
- [55] M. Grabisch and M. Roubens. An axiomatic approach to the concept of interaction among players in cooperative games. *International Journal of Game Theory*, 28(4):547–565, 1999.
- [56] J. Gu and V. Tresp. Contextual prediction difference analysis. *arXiv preprint arXiv:1910.09086*, 2019.
- [57] J. Gu, Y. Yang, and V. Tresp. Understanding individual decisions of cnns via contrastive backpropagation. In *14th Asian Conference on Computer Vision*, volume 11363 of *Lecture Notes in Computer Science*, pages 119–134. Springer, 2018.
- [58] R. Guidotti, A. Monreale, S. Ruggieri, D. Pedreschi, F. Turini, and F. Giannotti. Local rule-based explanations of black box decision systems. *arXiv preprint arXiv:1805.10820*, 2018.
- [59] R. Guidotti, A. Monreale, S. Ruggieri, F. Turini, F. Giannotti, and D. Pedreschi. A survey of methods for explaining black box models. *ACM Computing Surveys*, 51(5):93:1–93:42, 2019.
- [60] A. Gupta, J. Johnson, L. Fei-Fei, S. Savarese, and A. Alahi. Social GAN: socially acceptable trajectories with generative adversarial networks. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 2255–2264, 2018.
- [61] M. Hägele, P. Seegerer, S. Lapuschkin, M. Bockmayr, W. Samek, F. Klauschen, K.-R. Müller, and A. Binder. Resolving challenges in deep learning-based analyses of histopathological images using explanation methods. *Scientific Reports*, 10(1):6423, 2020.
- [62] K. Hansen, D. Baehrens, T. Schroeter, M. Rupp, and K.-R. Müller. Visual interpretation of kernel-based prediction models. *Molecular Informatics*, 30(9):817–826, 2011.

- [63] S. Haufe, F. C. Meinecke, K. Görgen, S. Dähne, J. Haynes, B. Blankertz, and F. Bießmann. On the interpretation of weight vectors of linear models in multivariate neuroimaging. *NeuroImage*, 87:96–110, 2014.
- [64] D. M. Hawkins. The detection of errors in multivariate data using principal components. *Journal of the American Statistical Association*, 69(346):340–344, 1974.
- [65] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016.
- [66] X. He, L. Liao, H. Zhang, L. Nie, X. Hu, and T. Chua. Neural collaborative filtering. In *Proceedings of the 26th International Conference on World Wide Web*, pages 173–182. ACM, 2017.
- [67] L. A. Hendricks, Z. Akata, M. Rohrbach, J. Donahue, B. Schiele, and T. Darrell. Generating visual explanations. In *14th European Conference on Computer Vision*, volume 9908 of *Lecture Notes in Computer Science*, pages 3–19. Springer, 2016.
- [68] J. Heo, S. Joo, and T. Moon. Fooling neural network interpretations via adversarial model manipulation. In *Advances in Neural Information Processing Systems 32*, pages 2921–2932, 2019.
- [69] M. A. Hernán and S. R. Cole. Invited commentary: causal diagrams and measurement bias. *American journal of epidemiology*, 170(8):959–962, 2009.
- [70] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997.
- [71] J. Hochuli, A. Helbling, T. Skaist, M. Ragoza, and D. R. Koes. Visualizing convolutional neural network protein-ligand scoring. *Journal of Molecular Graphics and Modelling*, 2018.
- [72] H. Hoffmann. Kernel PCA for novelty detection. *Pattern Recognition*, 40(3):863–874, 2007.
- [73] A. Holzinger. From machine learning to explainable ai. In *2018 World Symposium on Digital Intelligence for Systems and Machines (DISA)*, pages 55–66, 2018.
- [74] A. Holzinger, G. Langs, H. Denk, K. Zatloukal, and H. Müller. Causability and explainability of artificial intelligence in medicine. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 9(4):e1312, 2019.
- [75] S. Hong, D. Yang, J. Choi, and H. Lee. Interpretable text-to-image synthesis with hierarchical semantic layout generation. In *Explainable AI: Interpreting, Explaining and Visualizing Deep Learning*, volume 11700 of *Lecture Notes in Computer Science*, pages 77–95. Springer, 2019.
- [76] S. Hooker, D. Erhan, P.-J. Kindermans, and B. Kim. A benchmark for interpretability methods in deep neural networks. In *Advances in Neural Information Processing Systems*, pages 9734–9745, 2019.
- [77] F. Horst, S. Lapuschkin, W. Samek, K.-R. Müller, and W. I. Schödlhorn. Explaining the unique nature of individual gait patterns with deep learning. *Scientific Reports*, 9(2391), 2019.
- [78] G. B. Huang, M. Ramesh, T. Berg, and E. Learned-Miller. Labeled faces in the wild: A database for studying face recognition in unconstrained environments. Technical Report 07-49, University of Massachusetts, Amherst, 2007.
- [79] F. N. Iandola, M. W. Moskewicz, K. Ashraf, S. Han, W. J. Dally, and K. Keutzer. Squeezenet: Alexnet-level accuracy with 50x fewer parameters and <1mb model size. *arXiv preprint arXiv:1602.07360*, 2016.
- [80] B. K. Iwana, R. Kuroki, and S. Uchida. Explaining convolutional neural networks using softmax gradient layer-wise relevance propagation. In *IEEE International Conference on Computer Vision Workshops*, pages 4176–4185. IEEE, 2019.
- [81] J. D. Janizek, P. Sturmshels, and S. Lee. Explaining explanations: Axiomatic feature interactions for deep networks. *arXiv preprint arXiv:2002.04138*, 2020.
- [82] M. H. Jarrahi. Artificial intelligence and the future of work: Human-AI symbiosis in organizational decision making. *Business Horizons*, 61(4):577–586, 2018.
- [83] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional architecture for fast feature embedding. In *Proceedings of the 22nd ACM international conference on Multimedia*, pages 675–678, 2014.
- [84] G. Katz, C. W. Barrett, D. L. Dill, K. Julian, and M. J. Kochenderfer. Reluplex: An efficient SMT solver for verifying deep neural networks. In *CAV (1)*, volume 10426 of *Lecture Notes in Computer Science*, pages 97–117. Springer, 2017.
- [85] J. Kauffmann, M. Esders, G. Montavon, W. Samek, and K.-R. Müller. From clustering to cluster explanations via neural networks. *arXiv preprint arXiv:1906.07633*, 2019.
- [86] J. Kauffmann, K.-R. Müller, and G. Montavon. Towards explaining anomalies: A deep Taylor decomposition of one-class models. *Pattern Recognition*, 101:107198, 2020.
- [87] D. R. Kelley, Y. Reshef, M. Bileschi, D. Belanger, C. Y. McLean, and J. Snoek. Sequential regulatory activity prediction across chromosomes with convolutional neural networks. *Genome Research*, 28(5):739–750, 2018.
- [88] J. Khan, J. S. Wei, M. Ringnér, L. H. Saal, M. Ladanyi, F. Westermann, F. Berthold, M. Schwab, C. R. Antonescu, C. Peterson, and P. S. Meltzer. Classification and diagnostic prediction of cancers using gene expression profiling and artificial neural networks. *Nature Medicine*, 7(6):673–679, 2001.
- [89] B. Kim, M. Wattenberg, J. Gilmer, C. J. Cai, J. Wexler, F. B. Viégas, and R. Sayres. Interpretability beyond feature attribution: Quantitative testing with concept activation vectors (TCAV). In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 2673–2682. PMLR, 2018.
- [90] P.-J. Kindermans, K. T. Schütt, M. Alber, K.-R. Müller, D. Erhan, B. Kim, and S. Dähne. Learning how to explain neural networks: PatternNet and PatternAttribution. In *6th International Conference on Learning Representations*, 2018.
- [91] T. N. Kipf and M. Welling. Semi-supervised classification with graph convolutional networks. In *5th International Conference on Learning Representations*, 2017.
- [92] C. Kittel. *Introduction to Solid State Physics*, volume 8. Wiley New York, 2004.
- [93] F. Klauschen, K.-R. Müller, A. Binder, M. Bockmayr, M. Hägele, P. Seegerer, S. Wienert, G. Pruneri, S. de Maria, S. Badve, et al. Scoring of tumor-infiltrating lymphocytes: From visual estimation to machine learning. *Seminars in Cancer Biology*, 52:151–157, 2018.
- [94] M. Kohlbrenner, A. Bauer, S. Nakajima, A. Binder, W. Samek, and S. Lapuschkin. Towards best practice in explaining neural network decisions with LRP. In *2020 International Joint Conference on Neural Networks*, pages 1–7. IEEE, 2020.
- [95] N. Kokhlikyan, V. Miglani, M. Martin, E. Wang, J. Reynolds, A. Melnikov, N. Lunova, and O. Reblitz-Richardson. Pytorch Captum. <https://github.com/pytorch/captum>, 2019.
- [96] K. Kourou, T. P. Exarchos, K. P. Exarchos, M. V. Karamouzis, and D. I. Fotiadis. Machine learning applications in cancer prognosis and prediction. *Computational and Structural Biotechnology Journal*, 13:8–17, 2015.
- [97] F. Kratzert, M. Herrnegger, D. Klotz, S. Hochreiter, and G. Klambauer. Neuralhydrology - interpreting lstms in hydrology. In *Explainable AI: Interpreting, Explaining and Visualizing Deep Learning*, volume 11700 of *Lecture Notes in Computer Science*, pages 347–362. Springer, 2019.
- [98] O. Z. Kraus, L. J. Ba, and B. J. Frey. Classifying and segmenting microscopy images with deep multiple instance learning. *Bioinformatics*, 32(12):52–59, 2016.
- [99] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, pages 1106–1114, 2012.
- [100] W. Landecker, M. D. Thomure, L. M. A. Bettencourt, M. Mitchell, G. T. Kenyon, and S. P. Brumby. Interpreting individual classifications of hierarchical networks. In *IEEE Symposium on Computational Intelligence and Data Mining*, pages 32–38, 2013.
- [101] S. Lapuschkin, A. Binder, G. Montavon, K.-R. Müller, and W. Samek. Analyzing classifiers: Fisher vectors and deep neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2912–2920, 2016.
- [102] S. Lapuschkin, A. Binder, G. Montavon, K.-R. Müller, and W. Samek. The layer-wise relevance propagation toolbox for artificial neural networks. *Journal of Machine Learning Research*, 17(114):1–5, 2016.
- [103] S. Lapuschkin, A. Binder, K.-R. Müller, and W. Samek. Understanding and comparing deep neural networks for age and gender classification. In *Proceedings of the IEEE International Conference on Computer Vision Workshops (ICCVW)*, pages 1629–38, 2017.
- [104] S. Lapuschkin, S. Wäldchen, A. Binder, G. Montavon, W. Samek, and K.-R. Müller. Unmasking Clever Hans predictors and assessing what machines really learn. *Nature Communications*, 10(1):1096, 2019.
- [105] H. Larochelle and G. E. Hinton. Learning to combine foveal glimpses with a third-order Boltzmann machine. In *Advances in Neural Information Processing Systems 23*, pages 1243–1251, 2010.
- [106] S. M. Lauritsen, M. Kristensen, M. V. Olsen, M. S. Larsen, K. M. Lauritsen, M. J. Jørgensen, J. Lange, and B. Thiesson. Explainable artificial intelligence model to predict acute critical illness from electronic health records. *Nature Communications*, 11(1), 2020.

- [107] Y. LeCun, Y. Bengio, and G. Hinton. Deep learning. *Nature*, 521(7553):436, 2015.
- [108] Y. A. LeCun, L. Bottou, G. B. Orr, and K.-R. Müller. Efficient backprop. In *Neural Networks: Tricks of the Trade*, pages 9–48. Springer, 2012.
- [109] Y. Lee, C. D. Kummerow, and I. Ebert-Uphoff. Applying machine learning methods to detect convection using goes-16 abi data. *Atmospheric Measurement Techniques Discussions*, 2020:1–28, 2020.
- [110] K. Leino, S. Sen, A. Datta, M. Fredrikson, and L. Li. Influence-directed explanations for deep convolutional networks. In *IEEE International Test Conference*, pages 1–8. IEEE, 2018.
- [111] M. Lin, Q. Chen, and S. Yan. Network in network. In *International Conference of Learning Representations (ICLR)*, 2014.
- [112] Z. C. Lipton. The mythos of model interpretability. *ACM Queue*, 16(3):30, 2018.
- [113] H. Liu, Q. Yin, and W. Y. Wang. Towards explainable NLP: A generative explanation framework for text classification. In *Proceedings of the 57th Conference of the Association for Computational Linguistics*, pages 5570–5581. Association for Computational Linguistics, 2019.
- [114] Y. Liu, K. Barr, and J. Reinitz. Fully interpretable deep learning model of transcriptional control. *Bioinformatics*, 36(Supplement_1):i499–i507, 2020.
- [115] S. M. Lundberg, G. Erion, H. Chen, A. DeGrave, J. M. Prutkin, B. Nair, R. Katz, J. Himmelfarb, N. Bansal, and S.-I. Lee. From local explanations to global understanding with explainable ai for trees. *Nature Machine Intelligence*, 2(1):2522–5839, 2020.
- [116] S. M. Lundberg and S. Lee. A unified approach to interpreting model predictions. In *Advances in Neural Information Processing Systems 30*, pages 4765–4774, 2017.
- [117] S. Ma, X. Song, and J. Huang. Supervised group lasso with applications to microarray data analysis. *BMC Bioinformatics*, 8, 2007.
- [118] J. MacDonald, S. Wäldchen, S. Hauch, and G. Kutyniok. A rate-distortion framework for explaining neural network decisions. *arXiv preprint arXiv:1905.11092*, 2019.
- [119] J. MacQueen. Some methods for classification and analysis of multivariate observations. In *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Statistics*, pages 281–297. Berkeley, Calif., 1967. University of California Press.
- [120] Y. Matsui and T. Matsui. Np-completeness for calculating power indices of weighted majority games. *Theoretical Computer Science*, 263(1-2):305–310, 2001.
- [121] A. McGovern, R. Lagerquist, D. J. Gagne, G. E. Jergensen, K. L. Elmore, C. R. Homeyer, and T. Smith. Making the black box more transparent: Understanding the physical implications of machine learning. *Bulletin of the American Meteorological Society*, 100(11):2175–2199, 2019.
- [122] R. Memisevic. Learning to relate images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(8):1829–1846, 2013.
- [123] T. Miller. Explanation in artificial intelligence: Insights from the social sciences. *Artificial Intelligence*, 267:1–38, 2019.
- [124] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. A. Riedmiller, A. Fidjeland, G. Ostrovski, S. Petersen, C. Beattie, A. Sadik, I. Antonoglou, H. King, D. Kumaran, D. Wierstra, S. Legg, and D. Hassabis. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015.
- [125] G. Montavon. Gradient-based vs. propagation-based explanations: An axiomatic comparison. In *Explainable AI: Interpreting, Explaining and Visualizing Deep Learning*, volume 11700 of *Lecture Notes in Computer Science*, pages 253–265. Springer, 2019.
- [126] G. Montavon, A. Binder, S. Lapuschkin, W. Samek, and K.-R. Müller. Layer-wise relevance propagation: An overview. In *Explainable AI: Interpreting, Explaining and Visualizing Deep Learning*, volume 11700 of *Lecture Notes in Computer Science*, pages 193–209. Springer, 2019.
- [127] G. Montavon, S. Lapuschkin, A. Binder, W. Samek, and K.-R. Müller. Explaining nonlinear classification decisions with deep Taylor decomposition. *Pattern Recognition*, 65:211–222, 2017.
- [128] G. Montavon, W. Samek, and K.-R. Müller. Methods for interpreting and understanding deep neural networks. *Digital Signal Processing*, 73:1–15, 2018.
- [129] G. F. Montúfar, R. Pascanu, K. Cho, and Y. Bengio. On the number of linear regions of deep neural networks. In *Advances in Neural Information Processing Systems 27*, pages 2924–2932, 2014.
- [130] N. Morch, U. Kjems, L. K. Hansen, C. Svarer, I. Law, B. Lautrup, S. Strother, and K. Rehm. Visualization of neural networks using saliency maps. In *Proceedings of ICNN’95-International Conference on Neural Networks*, volume 4, pages 2085–2090, 1995.
- [131] A. Mordvintsev, C. Olah, and M. Tyka. Inceptionism: Going deeper into neural networks. *Google Research Blog*, 2015.
- [132] K.-R. Müller, S. Mika, G. Rätsch, K. Tsuda, and B. Schölkopf. An introduction to kernel-based learning algorithms. *IEEE Transactions on Neural Networks*, 12(2):181–201, 2001.
- [133] M. Narayanan, E. Chen, J. He, B. Kim, S. Gershman, and F. Doshi-Velez. How do humans understand explanations from machine learning systems? an evaluation of the human-interpretability of explanation. *arXiv preprint arXiv:1802.00682*, 2018.
- [134] A. M. Nguyen, A. Dosovitskiy, J. Yosinski, T. Brox, and J. Clune. Synthesizing the preferred inputs for neurons in neural networks via deep generator networks. In *Advances in Neural Information Processing Systems 29*, pages 3387–3395, 2016.
- [135] A. M. Nguyen, J. Yosinski, and J. Clune. Deep neural networks are easily fooled: High confidence predictions for unrecognizable images. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 427–436, 2015.
- [136] A. M. Nguyen, J. Yosinski, and J. Clune. Multifaceted feature visualization: Uncovering the different types of features learned by each neuron in deep neural networks. *arXiv preprint arXiv:1602.03616*, 2016.
- [137] R. Okuta, Y. Unno, D. Nishino, S. Hido, and C. Loomis. Cupy: A numpy-compatible library for nvidia gpu calculations. In *Advances in Neural Information Processing Systems*, 2017.
- [138] T. E. Oliphant. *A guide to NumPy*, volume 1. Trelgol Publishing USA, 2006.
- [139] G. Papadopoulos, P. J. Edwards, and A. F. Murray. Confidence estimation methods for neural networks: a practical comparison. *IEEE Transactions on Neural Networks*, 12(6):1278–1287, 2001.
- [140] Y. Park, B. Kwon, J. Heo, X. Hu, Y. Liu, and T. Moon. Estimating pm_{2.5} concentration of the conterminous united states via interpretable convolutional neural networks. *Environmental Pollution*, 256:113395, 2020.
- [141] E. Parzen. On estimation of a probability density function and mode. *The Annals of Mathematical Statistics*, 33(3):1065–1076, 1962.
- [142] V. Petsiuk, A. Das, and K. Saenko. RISE: randomized input sampling for explanation of black-box models. In *British Machine Vision Conference*, page 151, 2018.
- [143] B. Poulin, R. Eisner, D. Szafron, P. Lu, R. Greiner, D. S. Wishart, A. Fyshe, B. Pearcy, C. Macdonell, and J. Anvik. Visual explanation of evidence with additive classifiers. In *Proceedings, the 21st National Conference on Artificial Intelligence and the 18th Innovative Applications of Artificial Intelligence Conference*, pages 1822–1829, 2006.
- [144] K. Preuer, G. Klambauer, F. Rippmann, S. Hochreiter, and T. Unterthiner. Interpretable deep learning in drug discovery. In *Explainable AI: Interpreting, Explaining and Visualizing Deep Learning*, volume 11700 of *Lecture Notes in Computer Science*, pages 331–345. Springer, 2019.
- [145] G. Quellec, K. Charrière, Y. Boudi, B. Cochener, and M. Lamard. Deep image mining for diabetic retinopathy screening. *Medical Image Analysis*, 39:178–193, 2017.
- [146] S. Raghunath, A. E. U. Cerna, L. Jing, D. P. vanMaanen, J. Stough, D. N. Hartzel, J. B. Leader, H. L. Kirchner, M. C. Stumpe, A. Hafez, A. Nemani, T. Carbonati, K. W. Johnson, K. Young, C. W. Good, J. M. Pfeifer, A. A. Patel, B. P. Delisle, A. Alsaïd, D. Beer, C. M. Haggerty, and B. K. Fornwalt. Prediction of mortality from 12-lead electrocardiogram voltage data using a deep neural network. *Nature Medicine*, 26(6):886–891, 2020.
- [147] M. T. Ribeiro, S. Singh, and C. 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*, pages 1135–1144, 2016.
- [148] M. T. Ribeiro, S. Singh, and C. Guestrin. Anchors: High-precision model-agnostic explanations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 18, pages 1527–1535, 2018.
- [149] L. Rieger, P. Chormai, G. Montavon, L. Hansen, and K.-R. Müller. Structuring neural networks for more explainable predictions. In *Explainable and Interpretable Models in Computer Vision and Machine Learning*, pages 115–131. Springer, 2019.
- [150] R. Roscher, B. Bohn, M. F. Duarte, and J. Garcke. Explainable machine learning for scientific insights and discoveries. *IEEE Access*, 8:42200–42216, 2020.
- [151] F. Rosenblatt. The perceptron: A probabilistic model for information storage and organization in the brain. *Psychological Review*, 65(6):386–408, 1958.
- [152] A. S. Ross, M. C. Hughes, and F. Doshi-Velez. Right for the right reasons: Training differentiable models by constraining their

- explanations. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence*, pages 2662–2670, 2017.
- [153] R. Rothe, R. Timofte, and L. Van Gool. Dex: Deep expectation of apparent age from a single image. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 10–15, 2015.
- [154] C. Rudin. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5):206–215, 2019.
- [155] L. Ruff, J. R. Kauffmann, R. A. Vandermeulen, G. Montavon, W. Samek, M. Kloft, T. G. Dietterich, and K.-R. Müller. A unifying review of deep and shallow anomaly detection. *Proceedings of the IEEE*, pages 1–40, 2021.
- [156] L. Ruff, R. Vandermeulen, N. Goernitz, L. Deecke, S. A. Siddiqui, A. Binder, E. Müller, and M. Kloft. Deep one-class classification. In *International Conference on Machine Learning*, pages 4393–4402, 2018.
- [157] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. S. Bernstein, A. C. Berg, and F. Li. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, 115(3):211–252, 2015.
- [158] W. Samek, A. Binder, G. Montavon, S. Lapuschkin, and K.-R. Müller. Evaluating the visualization of what a deep neural network has learned. *IEEE Transactions on Neural Networks and Learning Systems*, 28(11):2660–2673, 2016.
- [159] W. Samek, G. Montavon, A. Vedaldi, L. K. Hansen, and K.-R. Müller, editors. *Explainable AI: Interpreting, Explaining and Visualizing Deep Learning*, volume 11700 of *Lecture Notes in Computer Science*. Springer, 2019.
- [160] F. Scarselli, M. Gori, A. C. Tsoi, M. Hagenbuchner, and G. Monfardini. The graph neural network model. *IEEE Transactions on Neural Networks*, 20(1):61–80, 2009.
- [161] J. Schmidhuber. Deep learning in neural networks: An overview. *Neural Networks*, 61:85–117, 2015.
- [162] T. Schnake, O. Eberle, J. Lederer, S. Nakajima, K. T. Schütt, K.-R. Müller, and G. Montavon. Higher-order explanations of graph neural networks via relevant walks. *arXiv preprint arXiv:2006.03589*, 2020.
- [163] B. Schölkopf and A. J. Smola. *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. Adaptive Computation and Machine Learning Series. MIT Press, 2002.
- [164] B. Schölkopf, A. J. Smola, and K.-R. Müller. Nonlinear component analysis as a kernel eigenvalue problem. *Neural Computation*, 10(5):1299–1319, 1998.
- [165] K. T. Schütt, F. Arbabzadah, S. Chmiela, K. R. Müller, and A. Tkatchenko. Quantum-chemical insights from deep tensor neural networks. *Nature Communications*, 8:13890, 2017.
- [166] K. T. Schütt, M. Gastegger, A. Tkatchenko, and K.-R. Müller. Quantum-chemical insights from interpretable atomistic neural networks. In *Explainable AI: Interpreting, Explaining and Visualizing Deep Learning*, volume 11700 of *Lecture Notes in Computer Science*, pages 311–330. Springer, 2019.
- [167] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra. Grad-CAM: Visual explanations from deep networks via gradient-based localization. In *IEEE International Conference on Computer Vision (ICCV)*, pages 618–626, 2017.
- [168] C. Shan. Learning local features for age estimation on real-life faces. In *Proceedings of the 1st ACM international workshop on Multimodal pervasive video analysis*, pages 23–28. ACM, 2010.
- [169] L. S. Shapley. 17. a value for n-person games. In *Contributions to the Theory of Games (AM-28), Volume II*. Princeton University Press, 1953.
- [170] A. Shrikumar, P. Greenside, and A. Kundaje. Learning important features through propagating activation differences. In *Proceedings of the 34th International Conference on Machine Learning (ICML)*, pages 3145–3153, 2017.
- [171] A. Shrikumar, P. Greenside, A. Shcherbina, and A. Kundaje. Not just a black box: Learning important features through propagating activation differences. *arXiv preprint arXiv:1605.01713*, 2016.
- [172] A. Shrikumar, J. Su, and A. Kundaje. Computationally efficient measures of internal neuron importance. *arXiv preprint arXiv:1807.09946*, 2018.
- [173] M. Simon, E. Rodner, T. Darrell, and J. Denzler. The whole is more than its parts? from explicit to implicit pose normalization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 42(3):749–763, 2020.
- [174] K. Simonyan, A. Vedaldi, and A. Zisserman. Deep inside convolutional networks: Visualising image classification models and saliency maps. In *International Conference of Learning Representations Workshops*, 2014.
- [175] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. In *3rd International Conference on Learning Representations*, 2015.
- [176] D. Smilkov, N. Thorat, B. Kim, F. B. Viégas, and M. Wattenberg. Smoothgrad: removing noise by adding noise. *arXiv preprint arXiv:1706.03825*, 2017.
- [177] C. Soneson, S. Gerster, and M. Delorenzi. Batch effect confounding leads to strong bias in performance estimates obtained by cross-validation. *PLoS ONE*, 9(6):e100335, 2014.
- [178] J. T. Springenberg, A. Dosovitskiy, T. Brox, and M. A. Riedmiller. Striving for simplicity: The all convolutional net. In *International Conference of Learning Representations (ICLR)*, 2015.
- [179] E. Strumbelj and I. Kononenko. An efficient explanation of individual classifications using game theory. *Journal of Machine Learning Research*, 11:1–18, 2010.
- [180] I. Sturm, S. Lapuschkin, W. Samek, and K.-R. Müller. Interpretable deep neural networks for single-trial eeg classification. *Journal of Neuroscience Methods*, 274:141–145, 2016.
- [181] P. Sturmfels, S. Lundberg, and S.-I. Lee. Visualizing the impact of feature attribution baselines. *Distill*, 5(1), Jan. 2020.
- [182] D. Su, H. Zhang, H. Chen, J. Yi, P. Chen, and Y. Gao. Is robustness the cost of accuracy? - A comprehensive study on the robustness of 18 deep image classification models. In *15th European Conference on Computer Vision*, volume 11216 of *Lecture Notes in Computer Science*, pages 644–661. Springer, 2018.
- [183] M. Sundararajan, A. Taly, and Q. Yan. Axiomatic attribution for deep networks. In *Proceedings of the 34th International Conference on Machine Learning*, pages 3319–3328, 2017.
- [184] W. R. Swartout and J. D. Moore. Explanation in second generation expert systems. In *Second Generation Expert Systems*, pages 543–585. Springer Berlin Heidelberg, 1993.
- [185] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. J. Goodfellow, and R. Fergus. Intriguing properties of neural networks. In *International Conference of Learning Representations (ICLR)*, 2014.
- [186] A. W. Thomas, H. R. Heekeren, K.-R. Müller, and W. Samek. Analyzing neuroimaging data through recurrent deep learning models. *Frontiers in Neuroscience*, 13:1321, 2019.
- [187] H. Trau Müller and A. Eriksson. The frequency range of the voice fundamental in the speech of male and female adults. *Unpublished manuscript*, 1995.
- [188] M. Tsang, D. Cheng, and Y. Liu. Detecting statistical interactions from neural network weights. In *6th International Conference on Learning Representations*, 2018.
- [189] B. Ustun, A. Spangher, and Y. Liu. Actionable recourse in linear classification. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, pages 10–19. ACM, 2019.
- [190] V. Vapnik. *The Nature of Statistical Learning Theory*. Springer, 1995.
- [191] M. M.-C. Vidovic, N. Görnitz, K.-R. Müller, G. Rätsch, and M. Kloft. Opening the black box: Revealing interpretable sequence motifs in kernel-based learning algorithms. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 137–153. Springer, 2015.
- [192] C. Von der Malsburg. Binding in models of perception and brain function. *Current Opinion in Neurobiology*, 5(4):520–526, 1995.
- [193] A. Warnecke, D. Arp, C. Wressnegger, and K. Rieck. Don’t paint it black: White-box explanations for deep learning in computer security. *arXiv preprint arXiv:1906.02108*, 2019.
- [194] C. K. Williams and C. E. Rasmussen. *Gaussian Processes for Machine Learning*. MIT press Cambridge, MA, 2006.
- [195] K. Xu, J. Ba, R. Kiros, K. Cho, A. C. Courville, R. Salakhutdinov, R. S. Zemel, and Y. Bengio. Show, attend and tell: Neural image caption generation with visual attention. In *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *JMLR Workshop and Conference Proceedings*, pages 2048–2057. JMLR.org, 2015.
- [196] Y. Yang, V. Tresp, M. Wunderle, and P. A. Fasching. Explaining therapy predictions with layer-wise relevance propagation in neural networks. In *IEEE International Conference on Healthcare Informatics*, pages 152–162, 2018.
- [197] I.-C. Yeh. Modeling of strength of high-performance concrete using artificial neural networks. *Cement and Concrete Research*, 28(12):1797–1808, 1998.
- [198] Z. Ying, D. Bourgeois, J. You, M. Zitnik, and J. Leskovec. Gnnexplainer: Generating explanations for graph neural networks. In

Advances in Neural Information Processing Systems 32, pages 9240–9251, 2019.

- [199] K. Young, G. Booth, B. Simpson, R. Dutton, and S. Shrapnel. Deep neural network or dermatologist? In *Interpretability of Machine Intelligence in Medical Image Computing and Multimodal Learning for Clinical Decision Support*, pages 48–55. Springer, 2019.
- [200] T. Zahavy, N. Ben-Zrihem, and S. Mannor. Graying the black box: Understanding dqns. In *Proceedings of the 33rd International Conference on Machine Learning*, volume 48 of *JMLR Workshop and Conference Proceedings*, pages 1899–1908. JMLR.org, 2016.
- [201] M. D. Zeiler and R. Fergus. Visualizing and understanding convolutional networks. In *European Conference Computer Vision - ECCV 2014*, pages 818–833, 2014.
- [202] J. Zhang, Z. Bargal, Sarah Adeland Lin, J. Brandt, X. Shen, and S. Sclaroff. Top-down neural attention by excitation backprop. *International Journal of Computer Vision*, 126(10):1084–1102, 2018.
- [203] Q. Zhang, X. Wang, R. Cao, Y. N. Wu, F. Shi, and S.-C. Zhu. Extracting an explanatory graph to interpret a cnn. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 1–1, 2020.
- [204] Z. Zhang, P. Chen, M. McGough, F. Xing, C. Wang, M. Bui, Y. Xie, M. Sapkota, L. Cui, J. Dhillon, et al. Pathologist-level interpretable whole-slide cancer diagnosis with deep learning. *Nature Machine Intelligence*, 1(5):236–245, 2019.
- [205] B. Zhou, D. Bau, A. Oliva, and A. Torralba. Interpreting deep visual representations via network dissection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(9):2131–2145, 2018.
- [206] B. Zhou, A. Khosla, À. Lapedriza, A. Oliva, and A. Torralba. Learning deep features for discriminative localization. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2921–2929, 2016.
- [207] X.-P. Zhu, J.-M. Dai, C.-J. Bian, Y. Chen, S. Chen, and C. Hu. Galaxy morphology classification with deep convolutional neural networks. *Astrophysics and Space Science*, 364(4):55, 2019.
- [208] L. M. Zintgraf, T. S. Cohen, T. Adel, and M. Welling. Visualizing deep neural network decisions: Prediction difference analysis. In *International Conference on Learning Representations (ICLR)*, 2017, 2017.



Wojciech Samek (M'13) is head of the Department of Artificial Intelligence and the Explainable AI Group at Fraunhofer Heinrich Hertz Institute, Berlin, Germany. He studied computer science at Humboldt University of Berlin, from 2004 to 2010, and received the Ph.D. degree with distinction from the Technical University of Berlin in 2014. During his studies he was awarded scholarships from the German Academic Scholarship Foundation and the DFG Research Training Group GRK 1589/1, and was a visiting researcher at NASA Ames Research

Center, Mountain View, USA. In 2014 he founded the Machine Learning Group at Fraunhofer HHI, which he has directed until 2020. Dr. Samek is associated faculty at the Berlin Institute for the Foundation of Learning and Data (BIFOLD), the ELLIS Unit Berlin and the DFG Graduate School BIOQIC. Furthermore, he is an editorial board member of Pattern Recognition, PLoS ONE and IEEE TNNLS, and an elected member of the IEEE MLSP Technical Committee. He has been serving as an AC for NAACL'21, was a recipient of multiple best paper awards, including the 2020 Pattern Recognition Best Paper Award, and a part of the MPEG-7 Part 17 standardization. His research interest include deep learning, explainable AI, neural network compression, and federated learning.



Grégoire Montavon is a senior researcher in the Machine Learning Group at the Technische Universität Berlin, and in the Berlin Institute for the Foundations of Learning and Data (BIFOLD). He received a Masters degree in Communication Systems from École Polytechnique Fédérale de Lausanne in 2009, and a Ph.D. degree in Machine Learning from the Technische Universität Berlin in 2013. He is a member of the ELLIS Unit Berlin, and an editorial board member of Pattern Recognition. He is recipient of the 2020 Pattern Recognition Best Paper Award. His research interests include explainable machine learning, deep neural networks, and unsupervised learning.



Sebastian Lapuschkin received the Ph.D. degree with distinction from the Berlin Institute of Technology in 2018 for his pioneering contributions to the field of eXplainable Artificial Intelligence and interpretable machine learning. From 2007 to 2013 he studied computer science (B.Sc. and M.Sc.) at the Berlin Institute of Technology, with a focus on software engineering and machine learning. Currently, he is a tenured researcher and head of the Explainable AI Group at Fraunhofer Heinrich Hertz Institute (HHI) in Berlin. He is recipient of multiple awards, including the Hugo-Geiger-Prize for outstanding doctoral achievement and the 2020 Pattern Recognition Best Paper Award. His research is focused on computer vision, (efficient) machine learning and data analysis, data and algorithm visualization, and the interpretation, holistic analysis and rectification of machine learning system behavior.



Christopher J. Anders is pursuing his Ph.D. as a junior researcher at Machine Learning Group at Technische Universität Berlin, and in the Berlin Institute for the Foundations of Learning and Data (BIFOLD). He received his B.Sc. in 2016 and his M.Sc. in 2018 both in computer science at Technische Universität Berlin, with a focus on machine learning and computer security. His research interests include (adversarial-/) explainable machine learning and deep neural networks.



Klaus-Robert Müller (M'12) has been a professor of computer science at Technische Universität Berlin since 2006; at the same time he is co-directing the Berlin Big Data Center. He studied physics in Karlsruhe from 1984 to 1989 and obtained his Ph.D. degree in computer science at Technische Universität Karlsruhe in 1992. After completing a postdoctoral position at GMD FIRST in Berlin, he was a research fellow at the University of Tokyo from 1994 to 1995. In 1995, he founded the Intelligent Data Analysis group at GMD-FIRST (later Fraunhofer FIRST) and directed it until 2008. From 1999 to 2006, he was a professor at the University of Potsdam. He was awarded the Olympus Prize for Pattern Recognition (1999), the SEL Alcatel Communication Award (2006), the Science Prize of Berlin by the Governing Mayor of Berlin (2014), the Vodafone Innovations Award (2017). In 2012, he was elected member of the German National Academy of Sciences-Leopoldina, in 2017 of the Berlin Brandenburg Academy of Sciences and also in 2017 external scientific member of the Max Planck Society. In 2019 and 2020 he became Highly Cited researcher in the cross-disciplinary area. His research interests are intelligent data analysis and Machine Learning in the sciences (Neuroscience (specifically Brain-Computer Interfaces), Physics, Chemistry).