



# HHS Public Access

Author manuscript

*IEEE Trans Radiat Plasma Med Sci.* Author manuscript; available in PMC 2022 November 01.

Published in final edited form as:

*IEEE Trans Radiat Plasma Med Sci.* 2021 November ; 5(6): 741–760. doi:10.1109/trpms.2021.3066428.

## On Interpretability of Artificial Neural Networks: A Survey

**Feng-Lei Fan [Student Member, IEEE],**

Department of Biomedical Engineering, Rensselaer Polytechnic Institute, Troy, NY, USA.

**Jinjun Xiong [Senior Member, IEEE],**

IBM Thomas J. Watson Research Center, Yorktown Heights, NY, 10598, USA.

**Mengzhou Li [Student Member, IEEE],**

Department of Biomedical Engineering, Rensselaer Polytechnic Institute, Troy, NY, USA.

**Ge Wang [Fellow, IEEE]**

Department of Biomedical Engineering, Rensselaer Polytechnic Institute, Troy, NY, USA.

### Abstract

Deep learning as represented by the artificial deep neural networks (DNNs) has achieved great success recently in many important areas that deal with text, images, videos, graphs, and so on. However, the black-box nature of DNNs has become one of the primary obstacles for their wide adoption in mission-critical applications such as medical diagnosis and therapy. Because of the huge potentials of deep learning, increasing the interpretability of deep neural networks has recently attracted much research attention. In this paper, we propose a simple but comprehensive taxonomy for interpretability, systematically review recent studies in improving interpretability of neural networks, describe applications of interpretability in medicine, and discuss possible future research directions of interpretability, such as in relation to fuzzy logic and brain science.

### Index Terms—

Deep learning; neural networks; interpretability; survey

## I. Introduction

Deep learning [71] has become the mainstream approach in many important domains targeting common objects such as text [40], images [182], videos [132], and graphs [88]. However, deep learning works as a black box model in the sense that, although deep learning performs quite well in practice, it is difficult to explain its underlying mechanism and behaviors. Questions are often asked such as how deep learning makes such a prediction, why some features are favored over others by a model, and what changes are needed to improve model performance, etc. Unfortunately, only modest success has been made to answer these questions.

---

Personal use is permitted, but republication/redistribution requires IEEE permission. See [http://www.ieee.org/publications\\_standards/publications/rights/index.html](http://www.ieee.org/publications_standards/publications/rights/index.html) for more information.

wangg6@rpi.edu .

Interpretability of deep neural networks is essential to many fields, and to healthcare [67], [68], [174] in particular for the following reasons. First, model robustness is a vital issue in medical applications. Recent studies suggest that model interpretability and robustness are closely connected [131]. On the one hand, the improvements in model robustness prompt model interpretability. For example, a deep model trained via adversarial training, a training method that augments training data with adversarial examples, shows better interpretability (with more accurate saliency maps) than the same model trained without adversarial examples [131]. On the other hand, when we understand a model deeply, we can thoroughly examine its weaknesses because the interpretability can help identify potential vulnerabilities of a complicated model, thereby improving its accuracy and reliability. Also, interpretability plays an important role in ethic use of deep learning techniques [57]. To build patients' trust in deep learning, interpretability is needed to hold a deep learning system accountable [57]. If a model builder can explain why a model makes a particular decision under certain conditions, users would know whether such a model contributes to an adverse event or not. It is then possible to establish standards and protocols to use the deep learning system optimally.

However, the lack of interpretability has become a main barrier of deep learning in its wide acceptance in mission-critical applications. For example, regulations were proposed by European Union in 2016 that individuals affected by algorithms have the right to obtain an explanation [61]. Despite great research efforts made on interpretability of deep learning and availability of several reviews on this topic, we believe that an up-to-date review is still needed, especially considering the rapid development of this area. The review of Q. Zhang and S. C. Zhu [202] is mainly on the visual interpretability. The representative publications from their review fall under the *feature analysis*, *saliency*, and *proxy* taxonomy in our review. The review of S. Chakraborty *et al.* [28] took opinions of [112] on levels of interpretability, and accordingly structured their review to provide in-depth perspectives but with limited scope. For example, only 49 references are cited there. The review of M. Du *et al.* [43] has a similar weakness, only covering 40 papers which are divided into post-hoc and ad-hoc explanations, as well as global and local interpretations. Their taxonomy is coarse-grained and neglects a number of important publications, such as publications on *explaining-by-text*, *explaining-by-case*, etc. In contrast, our review is much detailed and comprehensive, with the latest results included. While publications in L. H. Gilpin *et al.* [58] are classified into understanding the workflow of a neural network, understanding the representation of a neural network, and explanation producing, we cover all these aspects and also discuss the studies on how to prototype an interpretable neural network. Reviews by R. Guidotti *et al.* [65] and A. Adadi and M. Berrada [2] cover existing black-box machine learning models instead of focusing on neural networks. As a result, several hallmark papers on explaining neural networks are missing in their survey, such as the interpretation from the perspective of mathematics and physics.

A. B. Arrieta *et al.* [10] provides an extensive review on explainable AI (XAI), where concepts and taxonomies are clarified, and challenges are identified. While that review covers interpretability of AI/ML in general, our review is specific to deep neural networks and offers unique perspectives and insights. Specifically, our review is novel in the following senses: 1) We treat post-hoc and ad-hoc interpretability separately, because the former

explains the existing models, while the latter constructs interpretable ones; 2) we include widely-studied generative models, advanced mathematical/physical methods that summarize advances in deep learning theory, and the applications of interpretability in medicine; 3) important methods are illustrated with customized examples and publicly available codes through GitHub; and 4) interpretability research is a rapidly evolving field, and many research articles are published every year. Hence, our review should be a valuable and up-to-date addition to the literature.

Before we start our survey, let us first state three essential questions regarding interpretability: What does interpretability mean? Why is interpretability difficult? And how to build a good interpretation method? The first question has been well addressed in [112], and we include their statements here for completeness. The second question was partially touched in [112], [146], and we incorporate those comments and complement them with our own views. We provide our own perspectives on the third question.

### A. What Does Interpretability Mean?

Although the word “interpretability” is frequently used, people do not reach a consensus on the exact meanings of interpretability, which partially accounts for why current interpretation methods are so diverse. For example, some researchers explore post-hoc explanations for models, while some focus on the interplay mechanism between machineries of a model. Generally speaking, interpretability refers to the extent of human’s ability to understand and reason a model. Based on the categorization of [112], we summarize the implications of interpretability in different levels.

- **Simulatability**—Simulatability is considered as the understanding over the entire model. In a good sense, we can understand the mechanism of a model at the top level in a unified theoretical framework, one example is what was reported in [140]: a class of radial basis function (RBF) networks can be expressed by a solution to the interpolation problem with a regularization term, where a RBF network is an artificial neural network with RBFs as activation functions. In view of simulatability, the simpler the model is, the higher simulatability the model has. For example, a linear classifier or regressor is totally understandable. To enhance simulatability, we can change some facilities of models or use crafted regularization terms.
- **Decomposability**—Decomposability is to understand a model in terms of its components such as neurons, layers, blocks, and so on. Such a modularized analysis is quite popular in engineering fields. For instance, the inner working of a complicated system is factorized as a combination of functionalized modules. A myriad of engineering examples such as software development and optical system design have justified that a modularized analysis is effective. In machine learning, a decision tree is a kind of modularized methods, where each node has an explicit utility to judge if a discriminative condition is satisfied or not, each branch delivers an output of a judgement, and each leaf node represents the final decision after computing all attributes. Modularizing a neural network is advantageous to the optimization of the network design since we know the role of each and every component of the entire model.

- **Algorithmic Transparency**—Algorithmic Transparency is to understand the training process and dynamics of a model. The landscape of the objective function of a neural network is highly non-convex. The fact that deep models do not have a unique solution hurts the model transparency. Nevertheless, it is intriguing that current stochastic gradient descent (SGD)-based learning algorithms still perform efficiently and effectively. If we can understand why learning algorithms work, deep learning research and applications will be accelerated.

## B. Why Is Interpretability Difficult?

After we learn the meanings of interpretability, a question is what obstructs practitioners to obtain interpretability. This question was partially addressed in [146] in terms of *commercial barrier* and *data wildness*. Here, we complement their opinion with additional aspects on *human limitation* and *algorithmic complexity*. We believe that the hurdles to interpretable neural networks come from the following four aspects.

- **Human Limitation:** Expertise is often insufficient in many applications. Nowadays, deep learning has been extensively used in tackling intricate problems, which even professionals are unable to comprehend adequately. What's worse is that these problems are not uncommon. For example, in a recent study [46], we proposed to use an artificial neural network to predict pseudo-random events. Specifically, we fed 100,000 binary sequential digits into the network to predict the 100,001<sup>th</sup> digit in the sequence. In our prediction, the highly sophisticated hidden relationship was learned to beat a purely random guess with a  $3\sigma$  precision. Furthermore, it was conjectured that high sensitivity and efficiency of neural networks may help discriminate the fundamental differences between pseudo-randomness and real quantum randomness. In this case, it is no wonder that interpretability for neural networks will be missing, because even most talented physicists know little about the essence of this problem, let alone fully understand predictions of the neural network.
- **Commercial Barrier:** In the commercial world, there are strong motives for corporations to hide their models. First and foremost, companies profit from black-box models. It is not a common practice that a company makes capital out of totally transparent models [146]. Second, model opacity helps protect hard work from being reverse engineered. An effective black box is ideal in the sense that customers being served can obtain satisfactory results while competitors are not able to steal their intellectual properties easily [146]. Third, prototyping an interpretable model may cost too much in terms of financial, computational, and other resources. Existing open-sourced superior models are accessible to easily construct a well-performed algorithm for a specific task. However, generating reliable and consistent understanding to the behavior of the resultant model demands much more endeavors.
- **Data Wildness:** On the one hand, although it is a big data era, high quality data are often not accessible in many domains. For example, in the project of predicting electricity grid failure [146], the data base involves text documents, accounting data about electricity dating back to 1890s, and data from new manhole inspections. Highly heterogenous and inconsistent data hamper not only the accuracy of deep learning models but also the

construction of interpretability. On the other hand, real-world data have the character of high dimensionality, which suppresses reasoning. For example, given an MNIST image classification problem, the input image is of size  $28 \times 28 = 784$ . Hence the deep learning model tackling this problem has to learn an effective mapping of 784 variables to one of ten digits. If we consider the ImageNet dataset, the number of input variables goes up to  $512 \times 512 \times 3 = 768432$ .

- **Algorithmic Complexity:** Deep learning is a kind of large-scale, highly nonlinear algorithms. Convolution, pooling, non-linear activation, shortcuts, and so on contribute to variability of neural networks. The number of trainable parameters of a deep model can be on the order of hundreds million or even more. Despite that nonlinearity may not necessarily result in opacity (for example, a decision tree model is not linear but interpretable), deep learning's series of nonlinear operations indeed prevent us from understanding its inner working. In addition, recursiveness is another source of difficulty. A typical example is the chaos behavior resultant from nonlinear recursiveness. It is well-known that even a simple recursive mathematical model can lead to intractable dynamics [107]. In [175], it was proved that there are chaotic behaviors such as bifurcations even in simple neural networks. In chaotic systems, tiny changes of initial inputs may lead to huge outcome differences, adding to the complexity of interpretation methods.

### C. How to Build a Good Interpretation Method?

The third major issue is the criteria for assessing quality of a proposed interpretability method. Because existing evaluation methods are still premature, we propose five general and well-defined rules-of-thumb: *exactness*, *consistency*, *completeness*, *universality*, and *reward*. Our rules-of-thumb are fine-grained and focus on the characteristics of interpretation methods, compared to that described in [42]: application-grounded, human-grounded, and function-grounded.

- **Exactness:** Exactness means how accurate an interpretation method is. Is it just limited to a qualitative description or with a quantitative analysis? Generally, quantitative interpretation methods are more desirable than qualitative counterparts.
- **Consistency:** Consistency suggests that there is not any contradiction in an explanation. For multiple similar samples, a fair interpretation should produce consistent answers. In addition, an interpretation method should conform to the predictions of the authentic model. For example, the proxy-based methods are evaluated based on how closely they replicate the original golden model.
- **Completeness:** Mathematically, a neural network is to learn a mapping that best fits data. A good interpretation method should show effectiveness in support of the maximal number of data instances and data types.
- **Universality:** With the rapid development of deep learning, the deep learning armory has been substantially enriched. Such diverse deep learning models play important roles in a wide spectrum of applications. A driving question is whether we can develop a universal

interpreter that deciphers as many models as possible so as to save labor and time. But this is technically challenging due to the high variability among models.

- **Reward:** What are gains from the improved understanding of neural networks? In addition to the trust from practitioners and users, fruits of interpretability can be insights into network design, training, etc. Due to its black-box nature, using neural networks is largely a trial-and-error process with sometimes contradictory intuitions. A thorough understanding of deep learning will be instrumental to the research and applications of neural networks.

Briefly, our contributions in this review are three-folds: 1) We propose a comprehensive taxonomy for interpretability of neural networks and describe key methods with our insights; 2) we systematically illustrate interpretability methods as educational aids, as shown in Figures 3, 5, 6, 7, 9, 10, 16, 17; and 3) we shed light on future directions of interpretability research in terms of the convergence of neural networks and rule systems, the synergy between neural networks and brain science, and interpretability in medicine.

## II. A SURVEY ON INTERPRETATION METHODS

In this section, we first present our taxonomy and then review interpretability results under each category of our taxonomy. We enter the search terms “Deep Learning Interpretability”, “Neural Network Interpretability”, “Explainable Neural Network”, and “Explainable Deep Learning” into the Web of Science on Sep 22, 2020, with the time range from 2000 to 2019. The number of articles with respect to years is plotted in Figure 1, which clearly shows an exponential trend in this field. With the survey, our motive is to cover as many important papers as possible, and therefore we do not limit ourselves within Web of Science. We also search related articles using Google Scholar, PubMed, IEEE Xplore, and so on.

### A. Taxonomy Definition

As shown in Figure 2, our taxonomy is based on our surveyed papers and existing taxonomies. We first classify the surveyed papers into post-hoc interpretability analysis and ad-hoc interpretable modeling. Post-hoc interpretability analysis explains existing models and can be further classified into *feature analysis*, *model inspection*, *saliency*, *proxy*, *advanced mathematical/physical analysis*, *explaining-by-case*, and *explaining-by-text*, respectively. Ad-hoc interpretable modeling builds interpretable models and can be further categorized into *interpretable representation* and *model renovation*. In our proposed taxonomy, the class “*advanced mathematical/physical analysis*” is novel, but it is unfortunately missing in the previous reviews. We argue that this class is rather essential, because the incorporation of math/physics is critical in placing deep learning on a solid foundation for interpretability. In the following, we clarify the taxonomy definition and its illustration. We would like to underscore that one method may fall into different classes, depending on how one views it.

- **Post-hoc Interpretability Analysis**—Post-hoc interpretability is conducted after a model is well learned. A main advantage of post-hoc methods is that one does not need to compromise interpretability with the predictive performance since prediction and interpretation are two separate processes without mutual interference. However, a post-hoc

interpretation is usually not completely faithful to the original model. If an interpretation is 100% accurate compared to the original model, it becomes the original model. Therefore, any interpretation method in this category is more or less inaccurate. What is worse is that we often do not know the nuance [146]. Such a nuance makes it hard for practitioners to have a full trust to an interpretation method, because the correctness of the interpretation method is not guaranteed.

**Feature analysis**—Feature analysis techniques are centered in comparing, analyzing, and visualizing features of neurons and layers. Through feature analysis, sensitive features and ways to process them are identified such that the rationale of the model can be explained to some extent.

Feature analysis techniques can be applied to any neural networks and provide qualitative insights on what kinds of features are learned by a network. However, these techniques lack an in-depth, rigorous, and unified understanding, and therefore cannot be used to revise a model towards a higher interpretability.

**Model inspection**—Model inspection methods use external algorithms to delve into neural networks by systematically extracting important structural and parametric information on inner working mechanisms of neural networks.

Methods in this class are more technically accountable than those in *feature analysis* because analytical tools such as statistics are directly involved in the performance analysis. Therefore, the information gained by a model inspection method is more trustworthy and rewarding. In an exemplary study [184], finding important data routing paths is used as a way to understand the model. With such data routing paths, the model can be faithfully compressed to a compact one. In other words, interpretability improves the trustworthiness of model compression.

**Saliency**—Saliency methods identify which attributes of input data are most relevant to a prediction or a latent representation of a model. In this category, human inspection is involved to decide if a saliency map is plausible. A saliency map is useful. For example, if a polar bear always appears in a picture coupled with snow or ice, the model may have misused the information of snow or ice to detect the polar bear rather than real features of polar bears for detection. With a saliency map, this issue can be found and hence avoided.

Saliency methods are popular in interpretability research, however, extensive random tests reported that some saliency methods can be model independent and data independent [3], i.e., saliency maps offered by some methods can be highly similar to results produced with edge detectors. This is problematic because it means that those saliency methods fail to find the true attributes of the input that account for the prediction of the model. Consequently, a model-relevant and data-relevant saliency method should be developed in these cases.

**Proxy**—Proxy methods construct a simpler and more interpretable proxy that closely resembles a trained, large, complex, and black-box deep learning model. Proxy methods can be either local in a partial space or global in a whole solution space. The exemplary proxy

models include decision trees, rule systems, and so on. The weakness of proxy methods is the extra cost needed to construct a proxy model.

**Advanced mathematical/physical analysis**—Advanced mathematical/physical analysis methods put a neural network into a theoretical mathematics/physics framework, in which the mechanism of a neural network is understood with advanced mathematics/physics tools. This class covers theoretical advances of deep learning including non-convex optimization, representational power, and generalization ability.

A concern in this class is that, to establish a reasonable interpretation, unrealistic assumptions are sometimes made to facilitate a theoretical analysis, which may compromise the practical validity of the explanation.

**Explaining-by-case**—Explaining-by-case methods are along the line of case-based reasoning [90]. People favor examples. One may not be engaged by boring statistic numbers of a product but could be amazed while listening to other users' experience of using such a product. This philosophy wins the heart of many practitioners and intrigues the case-based interpretation for deep learning. Explaining-by-case methods provide representative examples that capture the essence of a model.

Methods in this class are interesting and inspiring. However, this practice is more like a sanity check instead of a general interpretation because not much information regarding the inner working of a neural network is understood from selected query cases.

**Explaining-by-text**—Explaining-by-text methods generate text descriptions in image-language joint tasks that are conducive to understanding the behavior of a model. This class can also include methods that generate symbols for explanation.

Methods in this class are particularly useful in image-language joint tasks such as generating a diagnostic report from an X-ray radiograph. However, explaining-by-text is not a general technique for any deep learning model because it can only work when a language module exists in a model.

- **Ad-hoc Interpretable Modeling**—Ad-hoc interpretable modeling eliminates the biases from the post-hoc interpretability analysis. Although it is generally believed that there is a trade-off between interpretability and model expressibility [123], it is still possible to find a model that is both powerful and interpretable. One notable example is the work reported in [30], where an interpretable two-layer additive risk model has won the first place in FICO Recognition Contest.

**Interpretable representation**—Interpretable representation methods employ regularization techniques to steer the optimization of a neural network towards a more interpretable representation. Properties such as decomposability, sparsity, and monotonicity can enhance interpretability. As a result, regularized features become a way to allow more interpretable models. Correspondingly, the loss function must contain a regularization term



for the purpose of interpretability, which restricts the original model to perform its full learning task.

**Model renovation**—Model renovation methods seek interpretability by the means of designing and deploying more interpretable machineries into a network. Those machineries include a neuron with purposely designed activation function, an inserted layer with a special functionality, a modularized architecture, and so on. The future direction is to use more and more explainable components that can at the same time achieve similar state-of-the-art performance for diverse tasks.

## B. Post-hoc Interpretability Analysis

- **Feature Analysis**—Inverting-based methods [41], [117], [164], [201] crack the representation of a neural network by inverting feature maps into a synthesized image. For example, A. Mahendran and A. Vedaldi [117] assumed that a representation of a neural network  $\Omega_0$  for an input image  $x_0$  was modeled as  $\Omega_0 = \Omega(x_0)$ , where  $\Omega$  is the neural network mapping, usually not invertible. Then, the inverting problem was formulated as finding an image  $x^*$  whose neural network representation best matches  $\Omega_0$ , i.e.,  $\arg \min_x \|\Omega(x) - \Omega_0\|^2 + \lambda R(x)$ , where  $R(x)$  is a regularization term representing prior

knowledge about the input image. The goal is to reveal the lost information by comparing differences between the inverted image and the original one. A. Dosovitskiy *et al.* [41] directly trained a new network with features generated by the model of interest as the input and images as the label, to invert features of intermediate layers to images. It was found that contours and colors could still be reconstructed even from deeper layer features. M. D. Zeiler *et al.* [201] designed a deconvolution network consisting of unpooling, rectification, deconvolution operations, to pair with the original convolutional network so that features could be inverted without training. In the deconvolution network, an unpooling layer is realized by using locations of maxima; rectification is realized by setting negative values to zero; and deconvolution layers use transposed filters.

Activation maximization methods [45], [128], [129], [169] devote to synthesizing images that maximize the output of a neural network or neurons of interest. The resulting images are referred as “deep dreams” as these can be regarded as dream images of a neural network or a neuron.

In [16], [85], [108], [197], [211], it was pointed out that information about a deep model could be extracted from each neuron. J. Yosinski *et al.* [197] straightforwardly inspected the activation values of neurons in each layer with respect to different images or videos. They found that live activation values that change for different inputs are helpful to understand how a model work. Y. Li *et al.* [108] contrasted features generated by different initializations to investigate if a neural network learns a similar representation when randomly initialized. The receptive field (RF) is a spatial extent over which a neuron connects with an input volume [111]. To investigate the size and shape of RF of a given input for a neuron, B. Zhou *et al.* [211] presented a network dissection method that first selected  $K$  images with high activation values for neurons of interest and then constructed 5,000 occluded images for each of  $K$  images, and then fed them into a neural network to observe the changes in

activation values for a given unit. A large discrepancy signals an important patch. Finally, the occluded images that have large discrepancy were re-centered and averaged to generate an RF. This network dissection method has been scaled to generative networks [17]. In addition, D. Bau *et al.* [16] scaled up a low-resolution activation map of a given layer to the same size as the input, thresholded the map into a binary activation map, and then computed the overlapping area between the binary activation map and the ground-truth binary segmentation map as an interpretability measure. A. Karpathy *et al.* [85] defined the gate in LSTM [73] to be either left or right saturated depending on its activation value being either less than 0.1 or more than 0.9. In this regard, neurons that are often right saturated are interesting because this means that these neurons can remember their values over a long period. Q. Zhang *et al.* [203] dissected feature relations in a network with the premise that the feature map of a filter in each layer can be activated by part patterns in the earlier layer. They mined part patterns layer by layer, discovered activation peaks of part patterns from the feature map of each layer, and constructed an explanatory graph to describe the relations of hierarchical features, with each node representing a part pattern and the edge between neighboring layers representing a co-activation relation.

- **Model Inspection**—The empirical influence function is to measure the dependence of an estimator on a sample [99]. P. W. Koh and P. Liang [89] applied the concept of the influence function to address the following question: Given a prediction for one sample, do other samples in the dataset have positive effects or negative effects on that prediction? This analysis could also help identify mis-annotated labels and outliers existing in the data. As Figure 3 shows, given a LeNet-5 like network, two harmful images for a given image are identified by the influence function.

A. Bansal *et al.* [12], H. Lakkaraju *et al.* [97], and Q. Zhang *et al.* [204] worked on the detection of failures or biases in a neural network. For example, A. Bansal *et al.* [12] developed a model-agnostic algorithm to identify which instances a neural network is likely to fail to provide any prediction for. In such a scenario, the model would instead give a warning like “Do not trust these predictions” as an alert. Specifically, they annotated all failed images with a collection of binary attributes and clustered these images in the attribute space. As a result, each cluster indicates a failure mode. To recognize those mislabeled instances with high predictive scores in the dataset efficiently, H. Lakkaraju *et al.* [97] introduced two basic speculations: The first is that mislabeling an instance with high confidence is due to the systematic biases instead of random perturbation, while the second is that each failed example is representative and informative enough. Then, they clustered the images into several groups and designed a multi-armed bandit search strategy by taking each group as a bandit that plans which group should be queried and sampled in each step. To discover representation biases, Q. Zhang *et al.* [204] utilized ground-truth relationships among attributes according to human’s common knowledge (fire-hot vs ice-cold) to examine if a mined attribute relationship by a neural network well fits the ground truth.

Y. Wang *et al.* [184] demystified a network by identifying critical data routes. Specifically, a gate control binary vector  $\lambda_k \in \{0, 1\}^{n_k}$ , where  $n_k$  is the number of neurons in the  $k^{\text{th}}$  layer,

was multiplied to the output of the  $k^{\text{th}}$  layer, and the problem of finding control gate values is formulated as searching  $\lambda_1, \dots, \lambda_K$ :

$$\arg \min_{\lambda_1, \dots, \lambda_K} d(f_{\theta}(x), f_{\theta}(x; \lambda_1, \dots, \lambda_K)) + \gamma \sum_k \|\lambda_k\|_1,$$

where  $f_{\theta}$  is the mapping represented by a neural network parameterized by  $\theta$ ,  $f_{\theta}(x; \lambda_1, \dots, \lambda_K)$  is the mapping when control gates  $\lambda_1, \dots, \lambda_K$  are enforced,  $d(\cdot, \cdot)$  is a distance measure,  $\gamma$  is a constant controlling the trade-off between the loss and regularization, and  $\|\cdot\|_1$  is the  $l_1$  norm such that  $\lambda_k$  is sparse. The learned control gates could expose the important data processing paths of a model. B. Kim *et al.* [86] developed the concept activated vector (CAV) that can quantitatively measure the sensitivity of the concept  $C$  with respect to any layer of a model. First, a binary linear classifier  $h$  was trained to distinguish between layer activations stimulated by two sets of samples:  $\{f_l(x) : x \in P_C\}$  and  $\{f_l(x) : x \notin P_C\}$ , where  $f_l(x)$  is the layer activation at the  $l^{\text{th}}$  layer, and  $P_C$  denotes data embodying the concept  $C$ . Then, the CAV was defined as the normal unit vector  $v_C^l$  to a hyperplane of the linear classifier that separated samples with and without the defined concept. Finally,  $v_C^l$  was used to calculate the sensitivity for a concept  $C$  in the  $l^{\text{th}}$  layer as the directional derivatives:

$$\begin{aligned} S_{C, k, l} &= \lim_{\epsilon \rightarrow 0} \frac{h_{l, k}(f_l(x) + \epsilon v_C^l) - h_{l, k}(f_l(x))}{\epsilon} \\ &= \nabla h_{l, k}(f_l(x)) v_C^l, \end{aligned}$$

where  $h_{l, k}$  denotes the logits of the trained binary linear classifier for the output class  $k$ . J. You *et al.* [196] mapped a neural network into a relational graph, and then studied the relationship between the graph structures of neural networks and their predictive performance through massive experiments (transcribed a graph into a network and implemented the network on a dataset). They discovered that the predictive performance of a network was correlated with two graph measures: the clustering coefficient and the average path length.

- **Saliency**—There is a plethora of methods to obtain a saliency map. Partial dependence plot (PDP) and individual condition expectation (ICE) [53], [59], [74] are model-agnostic statistical tools to visualize the dependence between the responsible variables and the predictive variables. To compute the PDP, suppose there are  $p$  input dimensions and let  $S, C \subseteq \{1, 2, \dots, p\}$  be two complementary sets, where  $S$  is the set one will fix, and  $C$  is the set one will change. Then the PDP for  $x_S$  is defined by  $f_S = \int f(x_S, x_C) dx_C$ , where  $f$  is the model. Compared with PDP, the definition of ICE is straightforward. The ICE curve at  $x_S$  is obtained by fixing  $x_C$  and varying  $x_S$ . Figure 4 shows a simple example on how to compute PDP and ICE, respectively.

A simple approach is to study the change of prediction after removing one feature, also known as leave-one-out attribution [4], [83], [105], [143], [212]. For example, A. Kádár *et al.* [83] utilized this idea to define an omission score:  $1 - \cos(\mathbf{h}(S), \mathbf{h}(S_{-i}))$ , where

$\text{cosine}(\cdot, \cdot)$  is the cosine distance,  $\mathbf{h}$  is the representation for a sentence,  $S$  is the full sentence, and  $S_{\setminus i}$  is the sentence without the  $i^{\text{th}}$  word, and analyzed the importance of each word. P. Adler *et al.* [4] proposed to measure an indirect influence for correlated inputs. For example, in a house loan decision system, race should not be a factor for decision-making. However, solely removing the race factor is not sufficient to rule out the effect of race because some remaining factors such as “zipcode” are highly concerned with race.

Furthermore, Shapley value from cooperative game theory was used in [6], [27], [39], [113], [115]. Mathematically, Shapley value of a set function  $\hat{f}$  with respect to the feature  $i$  is defined as

$$\text{Shapley}_i(\hat{f}) = \sum_{S \subseteq P \setminus \{i\}} \frac{(N - |S| - 1)! |S|!}{N!} (\hat{f}(S \cup \{i\}) - \hat{f}(S)),$$

where  $|\cdot|$  is the size of a set,  $P$  is a total player set of  $N$  players, and the set function  $\hat{f}$  maps each subset  $S \subseteq P$  to a real number. Furthermore, the definition of Shapley value can be twisted to the neural network function  $f$  by replacing the features in the input that are not in  $S$  with the zero value. Motivated by reducing the prohibitive computational cost incurred by combinatorial explosion, M. Ancona *et al.* [6] proposed a novel and polynomial-time approximation for Shapley values, which basically computed the expectation of a random coalition rather than enumerated each and every coalition. Figure 5 shows a simple example of how Shapley values can be computed for a fully connected layer network trained on California Housing dataset which includes eight attributes such as house age and room number as the inputs and the house price as the label.

Instead of removing one or more features, researchers also resort to gradients. K. Simonyan *et al.* [157], D. Smilkov *et al.* [161], M. Sundararajan *et al.* [168] and S. Singla *et al.* [160] utilized the idea of gradients to probe the saliency of an input. K. Simonyan *et al.* [157] calculated the first-order Taylor expansion of the class score with respect to image pixels, by which the first-order coefficients produce a saliency map for a class. D. Smilkov *et al.* [161] demonstrated that gradients as a saliency map show a correlation between attributes and labels, however, typically gradients are rather noisy. To remove noise, they proposed “SmoothGrad” that adds noise into the input image multiple times and averages the resultant gradient maps:  $\widehat{M}_c(x) = \frac{1}{N} \sum_{n=1}^N M_c^{(n)}(x + N(0, \sigma^2))$ , where  $M_c^{(n)}$  is a gradient map for a class  $c$ , and  $N(0, \sigma^2)$  is the Gaussian noise with  $\sigma$  as the standard variance. Basically,  $\widehat{M}_c(x)$  is a smoothed version of a salient map. M. Sundararajan *et al.* [168] set two fundamental requirements for saliency methods: (sensitivity) if only one feature is different between the input and the baseline, and the outputs of the input and the baseline are different, then this very feature should be credited by a non-zero attribution; (implementation invariance) the attributions for the same feature in two functionally equivalent networks should be identical. Noticing that earlier gradient-based saliency methods fail the above two requirements, they put forth integrated gradients, which is formulated as  $(x_i - x'_i) \int_0^1 \frac{\partial F(\mathbf{x}' + \alpha(\mathbf{x} - \mathbf{x}'))}{\partial x_i} d\alpha$ , where  $F(\cdot)$  is a neural network mapping,  $\mathbf{x} = (x_1, x_2, \dots, x_N)$  is an input, and  $\mathbf{x}' = (x'_1, x'_2, \dots, x'_N)$  is

the baseline satisfying  $\left. \frac{\partial}{\partial \mathbf{x}} F(\mathbf{x}) \right|_{\mathbf{x} = \mathbf{x}'} = 0$ . In practice, the integral can be transformed into

a discrete summation  $\frac{(x_i - x'_i)}{M} \times \sum_{m=1}^M \frac{\partial F(\mathbf{x}' + \frac{m}{M}(\mathbf{x} - \mathbf{x}'))}{\partial x_i}$ , where  $M$  is the number of steps

in the approximation of the integral. S. Singla *et al.* [160] proposed to use second-order approximations of a Taylor expansion to produce a saliency map so as to consider feature dependencies.

S. Bach *et al.* [11] proposed layer-wise relevance propagation (LRP) to compute the relevance of one attribute to a prediction by assuming that a model representation  $f(\mathbf{x})$  can be expressed as the sum of pixel-wise relevance  $R_p^l$ , where  $\mathbf{x}$  is an input image,  $l$  is the index of the layer, and  $p$  is the index of the pixel of  $\mathbf{x}$ . Thus,  $f(\mathbf{x}) = \sum_p R_p^L$ , where  $L$  is the final layer and  $R_p^L = \frac{w_p x_p^{L-1}}{\sum_p w_p x_p^{L-1}} f(\mathbf{x})$ , where  $w_p$  is the weight between pixel  $p$  of the  $(L-1)^{th}$  layer and the final layer. Given a feed-forward neural network, the pixel-wise relevance score  $R_p^1$  of an input is derived by calculating  $R_p^l = \sum_j \frac{z_{pj}}{\sum_{p'} z_{p'j}} R_j^{l+1}$  backwards with  $z_{pj} = x_p^l w_{pj}^{(l,l+1)}$ , where  $w_{pj}^{(l,l+1)}$  is the weight between the pixel  $p$  of layer  $l$  and the pixel  $j$  of the  $(l+1)^{th}$  layer. Furthermore, L. Arras *et al.* [9] extended LRP to recurrent neural networks (RNNs) for sentiment analysis. G. Montavon *et al.* [125] employed the whole first-order term of deep Taylor decomposition to produce a saliency map instead of just gradients. Suppose  $\hat{\mathbf{x}}$  is a well-chosen root for the function by a model  $f(\mathbf{x})$ :  $f(\hat{\mathbf{x}}) = 0$ , because  $f(\mathbf{x})$  can be decomposed as  $f(\mathbf{x}) = f(\hat{\mathbf{x}}) + \left( \left. \frac{\partial f}{\partial \mathbf{x}} \right|_{\mathbf{x} = \hat{\mathbf{x}}} \right)^T \cdot (\mathbf{x} - \hat{\mathbf{x}}) + \epsilon = 0 + \sum_i \left. \frac{\partial f}{\partial x_i} \right|_{\mathbf{x} = \hat{\mathbf{x}}} (x_i - \hat{x}_i) + \epsilon$ , where  $\epsilon$  is high-order terms, the pixel relevance for the pixel  $i$  is expressed as  $R_i = \left. \frac{\partial f}{\partial x_i} \right|_{\mathbf{x} = \hat{\mathbf{x}}} (x_i - \hat{x}_i)$ .

Inspired by the fact that even though a neuron is not fired, it is still likely to reveal useful information, A. Shrikumar *et al.* [156] proposed *DeepLIFT* to compute the difference between the activation of each neuron and its reference, where the reference is the activation of that neuron when the network is provided a reference input, and then backpropagate the difference to the image space layer by layer as LRP does. C. Singh *et al.* [159] introduced contextual decomposition whose layer propagation formula is  $\beta_i = \mathbf{W} \beta_{i-1} + \frac{|\mathbf{W} \beta_{i-1}|}{|\mathbf{W} \beta_{i-1}| + |\mathbf{W} \gamma_{i-1}|} \cdot b$  and  $\gamma_i = \mathbf{W} \gamma_{i-1} + \frac{|\mathbf{W} \gamma_{i-1}|}{|\mathbf{W} \beta_{i-1}| + |\mathbf{W} \gamma_{i-1}|} \cdot b$ , where  $\mathbf{W}$  is the weight matrix between the  $i^{th}$  and  $(i-1)^{th}$  layers and  $b$  is the bias vector. The restricting condition is  $g_i(x) = \beta_i(x) + \gamma_i(x)$ , where  $g_i(x)$  is the output of  $i^{th}$  layer.  $\beta_i(x)$  is considered as the contextual contribution of the input and  $\gamma_i(x)$  implies contribution of the input to  $g_i(x)$  that is not included in  $\beta_i(x)$ .

Figure 6 showcases the evaluation of raw gradients, SmoothGrad, IntegratedGrad, and Deep Taylor methods with a LeNet-5-like network. Among them, IntegratedGrad and Deep Taylor methods perform superbly on five digits.

Mutual-information measure to quantify the association between inputs and latent representations of a deep model can also similarly work as the saliency [63], [149], [194].

In addition, there are other methods to obtain saliency maps as well. A. S. Ross *et al.* [145] defined a new loss term  $\sum_i \left( A_i \frac{\partial}{\partial x_i} \sum_{k=1}^K \log(\hat{y}_k) \right)^2$  for training, where  $i$  is an index of a pixel,  $A_i$  is the binary mask to be optimized,  $\hat{y}_k$  is the  $k^{\text{th}}$  digit of the label, and  $K$  is the number of class. This loss is to penalize the sharpness of gradients towards a clearer interpretation boundary. R. C. Fong and A. Vedaldi [52] explored to learn the smallest region to delete, which is to find the optimal  $m^*$ :

$$m^* = \arg \min_{m \in [0, 1]^n} \lambda \| \mathbf{1} - m \|_1 + f_c(x_0; m),$$

where  $m$  is the soft mask,  $f_c(x_0; m)$  represents the loss of the network for an image  $x_0$  with the soft mask, and  $n$  is the number of pixels. T. Lei *et al.* [102] utilized a generator to specify segments of an original text as so-called rationales, which fulfill two conditions: 1) rationales should be sufficient as a replacement for the initial text; 2) rationales should be short and coherent. Deriving rationales is actually equivalent to deriving a binary mask, which can be regarded as a saliency map. Based on the above two constraints, the penalty term for a mask is formulated as:

$$\Omega(z) = \lambda_1 \| z \|_1 + \lambda_2 \sum_t |z_t - z_{t-1}|,$$

where  $z = [z_1, z_2, \dots]$  is a mask, the first term penalizes the number of rationales, and the second term is for smoothness.

The class activation map method (CAM [210]) and its variant [151] utilized global average pooling before a fully connected layer to derive the discriminative area. Specifically, let  $f_k(x, y)$  represent the  $k^{\text{th}}$  feature map, for a given class  $c$ , the input to the softmax layer is  $\sum_k w_k^c \sum_{x, y} f_k(x, y)$ , where  $w_k^c$  is the weight vector connecting the  $k^{\text{th}}$  feature map and the class  $c$ . The discriminative area is obtained as  $\sum_k w_k^c f_k(x, y)$ , which directly implies the importance of the pixel at  $(x, y)$  for class  $c$ . What's more, some weakly supervised learning methods such as M. Oquab *et al.* [135] can obtain discriminative areas as well. Specifically, they trained a network only with object labels, however, when they rescaled the feature maps produced by the max-pooling layer, it was surprisingly found that these feature maps were consistent with the locations of objects in the input.

- **Proxy**—There are about three ways to prototype a proxy. The first one is direct extraction. The gist of direct extraction is to construct a new interpretable model such as a decision tree [92], [192] or a rule-based system directly from the trained model. As far as the rule extraction is concerned, both decompositional [152] and pedagogical methods [147], [173] can be used. Pedagogical approaches extract rules that enjoy a similar input-output relationship with that of a neural network. These rules do not correspond to the weights and structure of the network. For example, the Validity Interval Analysis (VIA) [118] extracts rules in the following form:

IF (input  $\in$  a hypercube), THEN class is  $C_j$ .

R. Setiono and H. Liu [152] clustered hidden unit activation values based on the proximity of activation values. Then, the activation values of each cluster were denoted by their average activation values, at the same time kept the accuracy of the neural network as intact as possible. Next, the input data with the same average hidden unit activation value were clustered together to obtain a complete set of rules. In Figure 7, we illustrate obtained rules from a one-hidden-layer network using R. Setiono and H. Liu's method over the Iris dataset. In a neural network for a binary classification problem, the decision boundaries divide the input space into two parts, corresponding to two classes respectively. The explanation system HYPINV developed in E. W. Saad *et al.* [147] computes for each and every decision boundary hyperplane a tangent vector. The sign of an inner product between an input instance and a tangent vector will imply the position of the input instance relative to the decision boundary. Based on such a fact, a rule system can be established.

Lastly, some specialized networks such as ANFIS [80] and RBF networks [126], straightforwardly correspond to fuzzy logic systems. For example, an RBF network is equivalent to a Takagi-Sugeno rule system [172] that comprises rules such as "if  $x \in \text{set } A$  and  $y \in \text{set } B$ , then  $z = f(x, y)$ " [136]. Fuzzy logic interpretation in [48] considers each neuron/filter in a network as a generalized fuzzy logic gate. In this view, a neural network is nothing but a deep fuzzy logic system. Specifically, they analyzed a new type of neural networks, called quadratic networks, in which all the neurons are quadratic neurons that replace the inner product with the quadratic operation [47]. Their interpretation generalized fuzzy logic gates implemented by quadratic neurons, and then computed the entropy based on spectral information of fuzzy operations in a network. It was suggested that such an entropy could have deep connections with properties of minima and the complexity of neural networks.

The second one is called knowledge distillation [23] as Figure 8 shows. Although knowledge distillation techniques are mostly used for model compression, their principles can also be used for interpretability. The motif of knowledge distillation is that cumbersome models can generate relatively accurate predictions, assigning probabilities to all the possible classes, known as soft labels, that are more informative than one-hot labels. For example, a horse is more likely to be classified as a dog instead of a mountain. But with one-hot labeling, both the dog class and mountain class have zero probability. It was shown in [23] that, by the means of matching the logits of the original model, the generalization ability of the original cumbersome model could be transferred into a simpler model. Along this direction, an interpretable proxy model such as a decision tree [38], [186], a decision set [98], a global additive model [171], and a simpler network [75] were developed. For example, S. Tan *et al.* [171] used soft labels to train a global additive model in the form  $h_0 + \sum_i h_i(x_i) + \sum_{i \neq j} h_{ij}(x_i, x_j) + \sum_{i \neq j} \sum_{j \neq k} h_{ijk}(x_i, x_j, x_k) + \dots$ , where  $\{h_i\}_i$  could work as a feature saliency directly.

The last one is to provide a local explainer as a proxy. Local explainer methods locally mimic the predictive behaviors of neural networks. The basic rationale is that when a neural network is inspected globally, it looks complex. However, if we tackle it locally, the picture becomes clearer.

One typical local explainer is Local Interpretable Model-agnostic Explanation (LIME) [141], which synthesizes a number of neighbor instances by randomly setting elements of that sample to zero and computing the corresponding outcomes. Then, a linear regressor is used to fit synthesized instances, where the coefficients of the linear model signify the contributions of features. As Figure 9 shows, the LIME method is applied to a breast cancer classification model to identify which attributes are contributing forces for the model's benign or malignant prediction.

Y. Zhang *et al.* [207] pointed out the lack of robustness in the LIME explanation, which originates from sampling variance, sensitivity to the choice of parameters, and variation across different data points. *Anchor* [142] is an improved extension of LIME, which is to find the most important segments of an input such that the variability of the rest segments does not matter. Mathematically, *Anchor* searches a set:  $A = \{z \mid f(z) = f(x), z \in x\}$ , where  $f(\cdot)$  is a black-box model,  $x$  is the input, and  $z$  is the part of  $x$ . Another proposal Local Rule-based Explanation (LORE) was from [64]. The LORE takes advantage of the genetic algorithm to generate the balanced neighbors instead of random neighbors, thereby yielding high-quality training data that alleviates sampling variance of LIME.

- **Advanced Mathematical/Physical Analysis**—Y. Lu *et al.* [114] showed that many residual networks can be explained as discretized numerical solutions of ordinary differential equations, i.e., the inner-working of a residual block in ResNet [69] can be modeled as  $u_{n+1} = u_n + f(u_n)$ , where  $u_n$  is the output of the  $n^{\text{th}}$  block, and  $f(u_n)$  is the block operation. It was noticed that  $u_{n+1} = u_n + f(u_n)$  is a one-step finite difference approximation of an ordinary differential equation  $\frac{du}{dt} = f(u)$ . This idea inspired the invention of ODE-Net [32]. As Figure 10 shows, the starting point and the dynamics are tuned by an ODE-Net to fit a spiral.

N. Lei *et al.* [101] constructed an elegant connection between the Wasserstein generative adversarial network (WGAN [8]) and optimal transportation theory. They concluded that with low dimensionality hypothesis and the intentionally designed distance function, a generator and a discriminator can exactly represent each other in a closed form. Therefore, the competition between a discriminator and a generator in WGAN in the training is unnecessary.

In [154], it was proposed that the learning of a neural network is to extract the most relevant information in the input random variable  $X$  that pertains to an output random variable  $Y$ . Naively, for a feedforward neural network, the following inequality of mutual information holds:

$$I(Y; X) \geq I(Y; h_j) \geq I(Y; h_i) \geq I(Y; \hat{Y}),$$

where  $I(\cdot; \cdot)$  denotes mutual information,  $h_i, h_j$  are outputs of hidden layers ( $i > j$  means that the  $i^{\text{th}}$  layer is deeper), and  $\hat{Y}$  is a final prediction. Furthermore, S. Yu and J. C. Principe [198] employed an information bottleneck theory to gauge the mutual information states of symmetric layers in a stacked autoencoder as shown in Figure 11:



$$I(X; X') \geq I(T_1; T'_1) \geq \dots \geq I(T_K; T'_K).$$

However, it is tricky to estimate the mutual information since the probabilistic distribution of data is usually unknown a priori.

S. Kolouri *et al.* [91] built an integral geometric explanation for neural networks with a generalized Radon transform. Let  $X$  be a random variable for the input, which conforms to the distribution  $p_X$ , then we can derive a probability distribution function for the output of a neural network  $f_\theta(X)$  parametrized with  $\theta$ :  $p_{f_\theta}(z) = \int_X p_X(x) \delta(z - f_\theta(x)) dx$ , which is the generalized Radon transform, and the hypersurface is  $H(t, \theta) = \{x \in X | f_\theta(x) = t\}$ . In this regard, the transform by a neural network is characterized by the twisted hypersurfaces. H. Huang [77] used the mean-field theory to characterize the mechanism of dimensionality reduction by a deep network that assumes weights in each layer and input data following a Gaussian distribution. In his study, the self-covariance matrix of the output of the  $l^{\text{th}}$

layer was computed as  $\mathcal{C}$ , then the intrinsic dimensionality was defined as  $D = \frac{(\sum_{i=1}^N \lambda_i)^2}{\sum_{i=1}^N \lambda_i^2}$ ,

where  $\lambda_j$  is the eigenvalue of  $\mathcal{C}$ , and  $N$  is the number of eigenvalues. The quantity  $D/N$  was investigated across layers to analyze how compact representation are learned across layers. J. C. Ye *et al.* [193] utilized a framelet theory and low-rank Hankel matrix to represent signals in terms of their local and non-local bases, corresponding to convolution and generalized pooling operations. However, in their study the network structure was simplified in concatenating two ReLU units into a linear unit such that the nonlinearity from ReLU units could be circumvented. As far as advanced physic models are concerned, P. Mehta and D. C. Schwab [121] built an exact mapping from the Kadanoff variational renormalized group [82] to the restricted Boltzmann Machine (RBM) [148]. This mapping is independent of forms of the energy functions and can be scaled to any RBM.

Theoretical neural network studies are essential to interpretability as well. Currently, theoretical foundations of deep learning are primarily from three perspectives: representation, optimization, and generalization.

**Representation:** Let us include two examples here. The first example is to explain why deep networks is superior to shallow ones. Recognizing success of deep networks, L. Szymanski and B. McCane [170], D. Rolnick and M. Tegmark [144], N. Cohen *et al.* [37], H. N. Mhaskar and T. Poggio [124], R. Eldan and O. Shamir [44], and S. Liang and R. Srikant [109] justified that a deep network is more expressive than a shallow one. The basic idea is to construct a special class of functions that can be efficiently represented by a deep network but hard to be approximated by a shallow one. The second example is to understand utilities of shortcut connections of deep networks. A. Veit *et al.* [178] showed that residual connections can render a neural network to manifest an ensemble-like behavior. Along this direction, it was reported in [110] that with shortcuts, a network can be super slim to allow for universal approximation.

**Optimization:** Generally, optimizing a deep network is a NP-hard non-convex problem. The pervasive existence of saddle points [56] leads to that even finding a local minimum is also NP-hard [5]. Of particular interest to us is why an over-parametrized network can still be optimized well because a deep network is a kind of over-parametrized networks. The character of an over-parameterized network is that the number of parameters in a network exceeds the number of data instances. M. Soltanolkotabi *et al.* [163] showed that when data are Gaussian distributed and activation functions of neurons are quadratic, the landscape of an over-parameterized one-hidden-layer network allows global optimum to be searched efficiently. Q. Nguyen and M. Hein [130] demonstrated that with respect to linearly separable data, under assumptions on the rank of weight matrices of a feedforward neural network, every critical point of a loss function is a global minimum. Furthermore, A. Jacot *et al.* [78] showed that when the number of neurons in each layer of a neural network goes infinitely large, the training only renders small changes for the network function. As a result, the training of the network turns into the kernel ridge regression.

**Generalization:** Conventional generalization theory is incompetent to explain why a deep network can generalize well despite that the number of parameters of a deep network is many more than the number of samples. Recently proposed generalization bounds [127] that rely on the norm of weight matrices partially solved this problem. However, these bounds have an abnormal dependence on data that more data lead to a larger generalization bound, which apparently contradicts the common sense. We prospect that more efforts are needed to resolve the generalization puzzle satisfactorily [18], [122].

- **Explaining-by-Case**—Basically, case-based explanations present a case that is believed by a neural network to be most similar to the query case needing an explanation. Finding a similar case for explanation and selecting a representative case from data as the prototype [19] are basically the same thing and just use different metrics for similarity. While prototype selection is to find a minimal subset of instances that can represent the whole dataset, case-based explanations use the similarity metric based on the closeness of representations of a neural network, thereby exposing the hidden representation information of the neural network. In this light, case-based explanations are also related to deep metric learning [150].

As shown in Figure 12, E. Wallace *et al.* [181] employed the k-nearest neighbor algorithm to obtain the most similar cases for the query case in the feature space and then computed the percentage of the nearest neighbors belonging to the expected class as a measure for interpretability, suggesting how much a prediction is supported by data. C. Chen *et al.* [31] constructed a model that could dissect images by finding prototypical parts. Specifically, the pipeline of the model splits into multiple channels after convolutional layers, in which the function of each channel is expected to learn a prototypical part of the input such as the head or body of a bird. The decision for an input image is made based on the similarity of features of channels.

S. Wachter *et al.* [180] offered a novel case-based explanation method by providing a counterfactual case, which is an imaginary case that is close to the query but has a different output from that of the query. Counterfactual explanation provides the so-

called “closest possible case” or the smallest change to yield a different outcome. For example, counterfactual explanations may produce the following statement: “If you have a good striker, your team would have won this soccer game.” Coincidentally, techniques to generate a counterfactual explanation have been developed for the purpose of “adversarial perturbation”, i.e., structural attack [191]. Essentially, finding a closest possible case  $x'$  to the input  $x$  is equivalent to finding the smallest perturbation to  $x$  such that the classification result changes. For example, the following optimization can be built:

$$\arg \min_{x'} \lambda (f(x') - y')^2 + d(x, x'),$$

where  $\lambda$  is a constant,  $y'$  is a different label, and  $d(\cdot, \cdot)$  is chosen to be the Manhattan distance in hope that the input be minimally perturbed. Y. Goyal *et al.* [62] explored an alternative way to derive a counterfactual visual explanation. Given an image  $I$  with a label  $c$ , since the counterfactual visual explanation represents the change for the input that can force the model to yield a different prediction class  $c'$ , they selected an image  $I'$  with a label  $c'$  and managed to recognize the spatial region in  $I$  and  $I'$  such that the replacement of the recognized region would alter the model prediction from  $c$  to  $c'$ .

- **Explaining-by-Text**—Neural image captioning uses a neural network to produce a natural language description for an image. Despite that neural image captioning is initially not for network interpretability, descriptive language about images can tell the information about how a neural network analyzes an image. One representative method is from [84] that combines a convolutional neural network and a bidirectional recurrent neural network to obtain a bimodal embedding. Due to the hypothesis that the two embeddings representing similar semantics across two modalities should share the nearby locations of two spaces, the objective function is defined as

$$S_{IT} = \sum_{i \in g_I} \max_{t \in g_T} v_i^T s_t,$$

where  $v_i$  is the  $i^{th}$  image fragment in the set  $g_I$  and  $s_t$  is the  $t^{th}$  word in a sentence  $g_T$ . Another representative method is the attention mechanism [137], [179], [189], [190], where deep features are to align the corresponding text descriptions by a recursive neural network such as LSTM [73]. An explanation for deep features is provided by the corresponding words in the text and attention maps, which reflect which parts of an image attract the attention of the neural network.

As shown in Figure 13, in the  $k^{th}$  attention module that takes  $y_0, y_1, \dots, y_n$  as input, suppose its output is  $t_k = \sum_i y_i s_{ki} \cdot s_{k0}, s_{k1}, \dots, s_{kn}$  together form an attention map for  $t_k$  with respect to the associated word. However, S. Jain and B. C. Wallace [79] argued that an attention map is not qualified to work as an explanation because they observed that the attention map was not correlated with other importance measures of features such as gradient-based measures, and the change of attention weights yielded no changes in prediction.

### C. Ad-hoc Interpretable Modeling

- **Interpretable Representation**—Traditionally, regularization techniques for deep learning are primarily designed to avoid overfitting. However, it is also feasible to devise regularization techniques to enhance an interpretable representation in terms of decomposability [33], [165], [182], [205], monotonicity [195], non-negativity [34], sparsity [167], human-in-the-loop prior [96], and so on.

For example, X. Chen *et al.* [33] invented *InfoGAN* which is a simple but effective way to learn an interpretable representation. Traditionally, a generative adversarial network (GAN) [60] imposes no restrictions on how a generator utilizes the noise. In contrast, *InfoGAN* maximizes the mutual information between the latent codes and observations, forcing each dimension of noise to encode a semantic concept. Particularly, the latent codes are made of discrete categorical codes and continuous style codes. As shown in Figure 14, two style codes control the localized part and the digit rotation respectively.

Incorporating monotonicity constraints [195] is also useful to enhance interpretability. A monotonical relationship means when the value of a specified attribute increases, the predictive value of a model either increases or decreases. Such a simplicity promotes interpretability as well. J. Chorowski and J. M. Zurada [34] imposed non-negativity to weights of neural networks and argued that it could improve interpretability because it eliminated the cancellation and aliasing effects among neurons. A. Subramanian *et al.* [167] employed a k-sparse autoencoder for word embedding to promote sparsity in the embedding and claimed that this enhanced interpretability because a sparse embedding reduced the overlap between words. I. Lage *et al.* [96] proposed a novel human-in-the-loop evaluation in selecting a model. Specifically, a diverse set of models were trained and sent to users for evaluation. Users were asked to predict what the label of a data point would be assigned by a model  $M$ . The shorter the response time was, the better a user understood the model. Then, the model with the lowest response time was chosen.

- **Model Renovation**—L. Chu *et al.* [35] proposed to use piecewise linear functions as activations for a neural network (PLNN), thereby the decision boundaries of PLNN could be explicitly defined and further a closed-form solution could be derived for predictions of a network. As Figure 15 shown, F. Fan *et al.* [49] proposed Soft-Autoencoder (Soft-AE) by using adaptable soft-thresholding units in encoding layers and linear units in decoding layers. Consequently, Soft-AE can be interpreted as a learned cascaded wavelet adaptation system.

L. Fan [50] explained a neural network as a generalized Hamming network, whose neurons compute the generalized Hamming distance:  $h(\mathbf{x}, \mathbf{w}) = \sum_{l=1}^L w_l + \sum_{l=1}^L x_l - 2\mathbf{x} \cdot \mathbf{w}$  for an input  $\mathbf{x} = (x_1, \dots, x_L)$  and a weight vector  $\mathbf{w} = (w_1, \dots, w_L)$ . The bias term in each neuron is specified as  $b = -\frac{1}{2}(\sum_{l=1}^L w_l + \sum_{l=1}^L x_l)$  so that each neuron is a generalized Hamming neuron. In this regard, the function of the batch normalization is demystified as making the bias suitable for computation of the generalized Hamming distance. C. C. J. Kuo *et al.* [95] proposed a transparent design for constructing a feedforward convolutional network without the need of backpropagation. Specifically, filters in convolutional layers were built

by selecting principal components of PCA for outputs of earlier pooling layers. A fully connected layer was constructed by treating it as a linear-squared regressor.

D. A. Melis and T. Jaakkola [123] claimed that a neural network model  $f$  is interpretable if it has the form that  $f(x) = g(\theta_1(x) h_1(x), \dots, \theta_k(x) h_k(x))$ , where  $h_i(x)$  is the prototypical concept from the input  $x$  and  $\theta_i(x)$  is the relevance associated with that concept,  $g$  is monotonic and completely additively separable. Such a model can learn interpretable basis concepts and facilitate saliency analysis. Similarly, J. Vaughan *et al.* [177] designed a network structure to compatibly learn the function formulated as  $f(x) = \mu + \gamma_1 h_1(\beta_1^T x) + \gamma_2 h_2(\beta_2^T x) + \dots + \gamma_K h_K(\beta_K^T x)$ , where  $\beta_k$  is the projection,  $h_k(\cdot)$  represents the nonlinear transformation,  $\mu$  is the bias, and  $\gamma_k$  is the weighting factor. Such a model is more interpretable than a general network, because the function of this model has simpler partial derivatives that can simplify saliency analysis, statistical analysis, and so on.

C. Li *et al.* [104] proposed deep supervision by using prior hierarchical tasks on features of intermediate layers. Specifically, we have a dataset  $\{(x, y_1, \dots, y_m)\}$ , where labels  $y_1, \dots, y_m$  are hierarchical that  $y_j, j < i$  is a strict necessary condition for the existence of  $y_i, i > 1$ . Such a scheme introduces a modularized idea that through supervision of a specific task for an intermediate layer, the learning of that layer is steered towards the pre-specified task, thereby gaining interpretability.

T. Wang [183] proposed to use an interpretable and insertable substitute on a subset of data which the complex black-box model overkills. In their work, a rule set was built as an interpretable model to make a decision on the input data first. Those inputs which a rule set was handicapped to classify were passed into the black-box model for decision making. The logic of this hybrid predictive system is that an interpretable model for regular cases without compromising accuracy, a complex black-box model for complicated cases.

C. Jiang *et al.* [81] proposed finite automata-recurrent neural network (FA-RNN) that can be directly transformed into the regular expressions such that a good interpretability is extracted. The roadmap is that the constructed FA-RNN can be approximated into finite automata, and further transformed into regular expressions because finite automata and a regular expression are mutually convertible. In analogy, a regular expression can also be decoded into an FA-RNN as an initialization. FA-RNN is a good example to manifest the synergy between a rule system and a neural network.

### III. INTERPRETABILITY IN MEDICINE

These days, reports are often seen in the news that deep learning-based algorithms outperform experts or classic algorithms in the field of medicine [153]. Indeed, given an adequate computational power and well-curated datasets, a properly designed model can deliver competitive performance in most well-defined pattern recognition tasks. However, due to the high stakes of medicine-concerned applications, it is not sufficient to have a deep learning model that produces correct answers without an explanation. In this section, we focus on several exemplary papers concerning applications of interpretability methods in

medicine, and we organize the articles of relevance in accordance with the aforementioned taxonomy.

#### • Post-hoc Interpretability Analysis

**Feature analysis**—P. Van Molle *et al.* [176] visualized convolutional neural networks to assist decision-making for skin lesion classification. In their work, feature activations generated from the last two convolutional layers were rescaled to the size of an input image as the activation maps. Where a map has high activations were inspected. The activation strengths across different border types, skin colors, skin types, etc. were compared. The activation map exposed a risk that some unexpected regions had uncommonly high activations.

D. Bychkov *et al.* [24] utilized a model that combines a VGG-16 network [158] and an LSTM network [73] to predict five-year survival of colorectal cancer based on digitized tumor tissue samples. In their work, an RGB pathological image was split into many tiles. A VGG-16 network extracted a high-dimensional feature vector from each tile, which was then fed into an LSTM network to predict five-year survival. They used t-SNE [116] to map features learned by VGG-16 into a two-dimensional space for visualization and found that different classes of features of VGG-16 were well separated.

**Saliency**—I. Sturm *et al.* [166] applied a deep network with LRP [11] for the single-trial EEG [22] classification. The network entails two linear mean pooling layers before being activated or normalized. The feature importance score is assigned by LRP (S. Bach *et al.*, 2015).

J. R. Zech *et al.* [200] developed a deep learning model for chest radiography to classify patients into having pneumonia or not. Through interpretability analysis by CAM [210], they reported the risk that a deep learning model could make an incorrect decision by capturing features irrelevant to diseases, such as metal tokens.

O. Oktay *et al.* [134] combined attention gates with the decoder part of U-Net to cope with interpatient variation in organs' shapes and sizes. The proposed model can improve model sensitivity and accuracy by inhibiting representations of irrelevant regions. Aided by attention gates, they found that the model gradually shifted its attention to regions of interest.

D. Ardila *et al.* [7] proposed a deep learning algorithm that considers a patient's current and previous CT volumes to predict the risk of lung cancer. They used the integrated gradient method [168] to derive saliency maps and invited experienced radiologists to examine the fidelity of these maps. It turned out that in all cases, the readers strongly agreed that the model indeed focused on the nodules.

H. Lee *et al.* [100] reported an attention-assisted deep learning system for detection and classification of acute intracranial haemorrhage, where an attention map identified a region relevant to the disease. They evaluated the localization accuracy of the attention maps by

computing the proportion of bleeding points overlapping with the attention maps. Overall, it was found that 78.1% bleeding points were detected in the attention maps.

W. Caicedo-Torres and J. Gutierrez [25] proposed a multi-scale deep convolutional neural network for the mortality prediction based on the measurement of 22 different items in ICU such as the sodium index, urine output, etc. In their work, three temporal scales were represented by stacking convolutional kernels of dimensions  $3 \times 1$ ,  $6 \times 1$ , and  $12 \times 1$ . The saliency map by DeepLIFT [156] was utilized for interpretability.

H. Guo *et al.* [66] introduced an effective dual-stream network that conjugates extracted features from ResNet [69] and clinical prior knowledge to predict the mortality risk of patients based on low-dose CT images. To further testify the effectiveness of the proposed model, they utilized t-SNE [116] to reduce the dimensionality of feature maps of malignant and benign samples and found that malignant and benign features were well separated. Also, they applied CAM [210] to reveal that the deceased subjects correctly classified by the model were prone to have strong activations.

**Proxy**—Z. Che *et al.* (2016) applied knowledge distillation into a deep model to learn a gradient boosting tree [106] (GBT), that provides not only robust prediction performance but also a good interpretability in the context of electronic health record prediction. Specifically, they trained three deep models respectively, and then used predictions of deep models as labels to train a GBT model. Experiments on a Pediatric ICU dataset were reported that the GBT model maintained the prediction performance of deep models in terms of mortality and ventilator-free days.

S. Pereira *et al.* [138] combined global and local interpretation efforts for brain tumor segmentation and penumbra estimation in stroke lesions, where the global interpretability was derived from mutual information to sense the dependence between an input sample and the prediction, while the local interpretability was cast by a variant of LIME [141].

**Explaining-by-Case**—N. C. F. Codella *et al.* [36] employed saliency and explaining-by-case methods to explain a dermoscopic image analysis network which was jointly trained by disease labels with a triple-let loss. Specifically, the interpretability was gained by the discovered neighbors and localized regions that were most relevant to the distance from queries and neighbors.

**Explaining-by-Text**—Z. Zhang *et al.* [208] proposed an all-in-one network that read pathology bladder cancer images, generated diagnostic reports, retrieved images according to symptomatic descriptions, and visualized attention maps. They designed an auxiliary attention sharpening module to improve the discriminability of attention maps. Pathologists' feedbacks suggested that the explanatory maps tended to highlight regions that concern with carcinoma-informative regions.

#### • Ad-hoc Interpretable Modeling

**Interpretable Representation**—X. Fang and P. Yan [51] devised the Pyramid Input Pyramid Output Feature Abstraction Network (PIPO-FAN) with multiple arms for multi-

organ segmentation. Each of the arm handles the information on one scale. The total loss is obtained by adding the segmentation loss to each of these arms such that segmentation-wise features are generated in each arm. Visualization analysis suggested that features from different arms have hierarchical semantical meanings, i.e., some are blurry but contain global class-wise information, while the others contain local boundary information. As shown in Figure 16, the segmentation loss creates semantically meaningful features, where low-scale arms produce more details and high-scale arms find global morphologies.

**Model Renovation**—W. Gale *et al.* [55] combined a DenseNet (G. Huang *et al.*, 2017) model with an LSTM model [73] for detection of hip features from pelvic X-ray radiographs. A radiologist hand-labelled standard descriptive terms to construct a semantic dataset for these radiographs. Their model consistently generated informative sentences favored by doctors over saliency maps. Also, they demonstrated that the combination of visualization and text interpretation give an interpretation superior to either of them alone.

C. Biffi *et al.* [20] employed a variational autoencoder [87] (VAE)-based model for classification of cardiac diseases as well as structurally remodeling based on cardiovascular images. In their scheme, registered left ventricular (LV) segmentations at end-diastolic (ED) and end-systolic (ES) phases were encoded in a low-dimensional latent space by VAE. The learned latent low dimensional manifold was connected to a multilayer perceptron (MLP) for disease classification. The interpretation was given by an activation maximization technique. The “deep dream” of MLP was derived and inverted to the image space for visualization.

S. Shen *et al.* [155] built an interpretable deep hierarchical semantic convolutional neural network (HSCNN) to predict the malignancy of pulmonary nodules in CT images. HSCNN consists of three modules: a general feature learning module, a low-level task module that predicts semantic characteristics such as sphericity, margin, subtlety, and so on, and a high-level task module absorbs information from both general features and low-level task predictions to produce an overall lung nodule malignancy. Due to the semantic meaning contained in the low-level task, HSCNN has boosted interpretability.

Z. Zhang *et al.* [209] developed a deep convolutional network to automate the whole-slide reading of pathology images for tumors and the diagnosis process of pathologists. Specially, the network can generate a clinical pathology report along with attention-assisted features.

Y. Lei *et al.* [103] observed that CAM [210] and Grad-CAM [151] are for interpreting localization tasks and tend to ignore fine-grained structures. Consequently, they proposed a shape-and-margin-aware soft activation map (SAM) that could probe subtle but critical features in a lung nodule classification task. The comprehensive experimental comparisons showed that compared to CAM and Grad-CAM, SAM can reveal relatively discrete and irregular features around nodules.

#### IV. PERSPECTIVE

In this section, we suggest a few directions, in hope to advance the understanding and practice of artificial neural networks.



## • Synergy of Fuzzy Logic and Deep Learning

Fuzzy logic [199] was a buzz phrase in the last nighties. It extends the Boolean logic from 0–1 judgement to imprecise inference with fuzziness in the interval [0, 1]. Fuzzy theory can be divided into two branches: fuzzy set theory and fuzzy logic theory. The latter, with an emphasis on “IF-THEN” rules, has demonstrated effectiveness in dealing with a plethora of complicated system modeling and control problems. Nevertheless, a fuzzy rule-based system is restricted by the acquisition of a large number of fuzzy rules, a process that is tedious and computationally expensive. While a neural network is a data-driven method that extracts knowledge from data through training, with the knowledge represented by neurons in a distributed manner. However, a neural network falls short of delivering a satisfactory result in the context of small data and suffers from the lack of interpretability. In contrast, a fuzzy logic system employs experts’ knowledge and represents a system in the form of IF-THEN rules. Although a fuzzy logic system merits interpretability and accountability, it is incompetent in efficient and effective knowledge acquisition. It seems that a neural network and a fuzzy logic system are complementary to each other. Therefore, it is instrumental to combine the best of two worlds towards an enhanced interpretability. In fact, this roadmap is not totally new. There have been several combinations along this direction: ANFIS model [80], generic fuzzy perceptron [126], RBF networks [21], and so on.

One suggestion is to build a deep RBF network. Given the input vector  $\mathbf{x} = [x_1, x_2, \dots, x_n]$ , an RBF network is expressed as  $f(\mathbf{x}) = \sum_i^n w_i \phi_i(\mathbf{x} - \mathbf{c}_i)$ , where  $\phi_i(\mathbf{x} - \mathbf{c}_i)$  is usually selected as  $\exp\left(-\frac{\|\mathbf{x} - \mathbf{c}_i\|^2}{2\sigma^2}\right)$ , where  $\mathbf{c}_i$  is the cluster center of the  $i^{\text{th}}$  neuron. It was proved the functional equivalence between an RBF network and a fuzzy inference system under mild conditions [21]. Also, an RBF network is shown to be a universal approximator [136]. Hence, an RBF network is a potentially sound vehicle that can encode fuzzy rules into its adaptive representation without loss of accuracy. Reciprocally, rule generation and fuzzy rule representation in an adaptable RBF network are more straightforward compared to a multilayer perceptron. Although current RBF networks are of one-hidden-layer structures, it is feasible to develop deep RBF networks, which can be viewed as a deep fuzzy rule system. A greedy layer-wise training algorithm was developed in [71], which successfully solved the training problem for deep networks. It is possible to translate such success into the training of deep RBF networks. Then, the correspondence between a deep RBF network and a deep fuzzy logic system will be applied to obtain a deep fuzzy rule system. We believe that efforts should be made to synergize fuzzy logic and deep learning techniques aided by big data along this direction.

## • Convergence of Neuroscience and Deep Learning

Up to date, truly intelligent systems are still only human. The artificial neural networks in their earlier forms were clearly inspired by biological neural networks [120]. However subsequent developments of neural networks were, to a much less degree, pushed by neurological and biological insights. As far as interpretability is concerned, since biological and artificial neural networks are deeply connected, advances in neuroscience should be relevant and even instrumental to the development and interpretation of deep learning

techniques. We believe that the neuroscience promises a bright future of deep learning interpretability in the following aspects.

**Cost function.**—The effective use of cost functions is a key driving force for the development of deep networks in the past years; for example, the adversarial loss used in GANs [60]. In previous sections, we have highlighted cases which demonstrate that an appropriate cost function will enable a model to learn an interpretable representation, such as enhance feature disentanglement. Along this direction, a myriad of cost functions can be built to reflect biologically plausible rationales. Indeed, our brain can be modeled as an optimization machine [119], which has a powerful credit assignment mechanism to form a cost function.

**Optimization algorithm.**—Despite the huge success achieved by backpropagation, it is far from ideal in the view of neuroscience. Truly in many senses, backpropagation fails to manifest the true behaviors of how a human neural system tunes the synapses of a neuron. For example, in a biological neural system, synapses are updated in a local manner [94] and only depend on the activities of presynaptic and postsynaptic neurons. However, connections in deep networks are tuned through non-local backpropagation. Figure 17 shows a bio-plausible learning algorithm for a two-layer network on CIFAR-100 [93]. Additionally, a neuromodulator is missing in deep networks in contrast to the inner-working of a human brain, where the state of one neuron can exhibit different input-output patterns controlled by a global neuromodulator like dopamine, serotonin, and so on [162]. Neuromodulators are believed to be critical due to their ability to selectively control on and off states of one neuron which is equivalently switching the involved cost function [13].

Considering that there are quite few studies discussing the interpretability of training algorithms, powerful and interpretable training algorithms will be highly desirable. Just like for classic optimization methods, we wish that future non-convex optimization algorithms will have some kinds of uniqueness, stability, and continuous dependency on data, etc.

**Bio-Plausible Architectural Design.**—In the past decades, neural networks were designed in diverse architectures from simple feedforward networks to deep convolutional networks and other highly sophisticated networks. The structure determines functionality, i.e., a specific network architecture regulates the information flow with distinct characteristics. Therefore, specialized architectures are useful as effective solutions for intended problems. Currently, the structural differences between deep learning and biological systems are eminent. A typical network is used and tuned for most tasks based on big data, while a biological system learns from a small number of data and generalizes very well. Clearly, a huge amount of knowledge needs to be learned from biological neural networks so that a more desirable and explainable neural network architectures can be designed.

#### • Interpretability in Medicine

A majority of interpretability research efforts in medicine are only for classification tasks, but radiological practices cover a large variety of tasks such as image segmentation,

registration, reconstruction, and so on. Clearly, interpretability is also closely relevant to these areas, and therefore it is in need to promote interpretability research in these domains. On the one hand, more efforts should be made to extend the existing interpretation methods to other tasks that have not been explored. On the other hand, practitioners can design task-specific interpretation methods with their expertise and insights. For example, in image segmentation, explaining why a voxel receives a class label in image segmentation is much harder than explaining which area in the input image is responsible for a prediction in image classification. Similarly, for image reconstruction, interpretability could be quite complicated. In this regard, our recently proposed ACID framework allows a synergistic integration of data-driven priors and compressed sensing (CS)-modeled priors, enforcing both of which iteratively via physics-based analytic mapping [188]. By doing so, modern CS and state-of-the-art deep networks are united to overcome the vulnerabilities of existing deep reconstruction networks, at the same time transferring the interpretability of the model-based methods to the hybrid deep neural networks.

In addition to the above referenced publications, gaining interpretability ultimately also relies on medical doctors, who have invaluable professional training despite some biases and errors. As a result, active collaboration among medical doctors, technical experts, and theoretical researchers to design effective, efficient, and reproducible ways to assess and apply interpretability methods will be an important avenue for future development of deep learning methods.

## V. CONCLUSION

In conclusion, we have reviewed key ideas, implications, limitations of existing interpretability studies, and illustrated some typical interpretation methods through examples. In doing so, we have depicted a holistic landscape of interpretability research using the proposed taxonomy and introduced applications of interpretability in medicine particularly. Figures 3, 5, 6, 7, 9, 10, 16, 17 are visualization results from our own implementation of chosen interpretation methods. We have open-sourced relevant codes in the GitHub (<https://github.com/FengleiFan/IndependentEvaluation>). There is no doubt that a unified and accountable interpretation framework is critical to elevate interpretability research into a new phase. In the future, more efforts are needed to reveal the essence of deep learning. Because this field is still highly interdisciplinary and rapidly evolving, there are great opportunities ahead that will be both academically and practically rewarding.

## ACKNOWLEDGEMENT

The authors are grateful for Dr. Hongming Shan's suggestions (Fudan University) and anonymous reviewers' advice.

This work was supported in part by the Rensselaer-IBM AI Research Collaboration Program (<http://airc.rpi.edu>), part of the IBM AI Horizons Network (<http://ibm.biz/AIHorizons>)

## References

- [1]. Aamodt A and Plaza E, "Case-based reasoning: Foundational issues, methodological variations, and system approaches," *AI communications*, vol. 7, no. 1, pp. 39–59, 1994.

- [2]. Adadi A and Berrada M, "Peeking inside the black-box: A survey on Explainable Artificial Intelligence (XAI)," *IEEE Access*, vol. 6, pp. 52138–52160, 2018.
- [3]. Adebayo J, Gilmer J, Muelly M, Goodfellow I, Hardt M, and Kim B, "Sanity checks for saliency maps," In *NeurIPS*, 2018.
- [4]. Adler P, Falk C, Friedler SA, Nix T, Rybeck G, Scheidegger C, Smith B, Venkatasubramanian S, "Auditing black-box models for indirect influence," *Knowledge and Information Systems*, vol. 54, no. 1, pp. 95–122, 2018.
- [5]. Anandkumar A and Ge R, "Efficient approaches for escaping higher order saddle points in non-convex optimization," In *COLT*, pp. 81–102, 2016.
- [6]. Ancona M, Öztireli C and Gross, "Explaining Deep Neural Networks with a Polynomial Time Algorithm for Shapley Values Approximation," In *ICML*, 2019.
- [7]. Ardila D, et al. , "End-to-end lung cancer screening with three-dimensional deep learning on low-dose chest computed tomography," *Nature medicine*, vol. 25, no. 6, pp. 954–961, 2019.
- [8]. Arjovsky M, Chintala S, Bottou L, "Wasserstein gan," *arXiv preprint arXiv:1701.07875*. 2017.
- [9]. Arras L, Montavon G, Müller KR and Samek W, "Explaining recurrent neural network predictions in sentiment analysis. *arXiv preprint arXiv:1706.07206*, 2017.
- [10]. Arrieta AB, Díaz-Rodríguez N, Del Ser J, et al. "Explainable Artificial Intelligence (XAI): Concepts, Taxonomies, Opportunities and Challenges toward Responsible AI," *Information Fusion*, 2019.
- [11]. Bach S, Binder A, Montavon G, Klauschen F, Müller KR & Samek W, "On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation," *PLoS one*, vol. 10, no. 7, e0130140, 2015. [PubMed: 26161953]
- [12]. Bansal A, Farhadi A and Parikh D, "Towards transparent systems: Semantic characterization of failure modes," In *ECCV*, 2014.
- [13]. Bargmann CI, "Beyond the connectome: how neuromodulators shape neural circuits," *Bioessays*, vol. 34, no. 6, pp. 458–65, 2012. [PubMed: 22396302]
- [14]. Bartlett PL, Foster DJ, and Telgarsky MJ, "Spectrally-normalized margin bounds for neural networks," In *NeurIPS*, pp. 6240–6249, 2017.
- [15]. Bastani O, Kim C, Bastani H, "Interpretability via model extraction," *arXiv preprint, arXiv:1706.09773*. 2017.
- [16]. Bau D, Zhou B, Khosla A, Oliva A, Torralba A, "Network dissection: Quantifying interpretability of deep visual representations," In *CVPR*, 2017.
- [17]. Bau D, Zhu JY, Strobelt H, Lapedriza A, Zhou B, Torralba A, "Understanding the role of individual units in a deep neural network," *Proceedings of the National Academy of Sciences*, 2020 Sep 1.
- [18]. Belkin M, Hsu D, Ma S, Mandal S, "Reconciling modern machine-learning practice and the classical bias–variance trade-off," *Proceedings of the National Academy of Sciences*, vol. 116, no. 32, pp. 15849–54, 2019.
- [19]. Bien J, Tibshirani R, "Prototype selection for interpretable classification," *The Annals of Applied Statistics*, vol. 5, no. 4, pp. 2403–24, 2011.
- [20]. Biffi C, Oktay O, Tarroni G, Bai W, De Marvaio A, Doumou G, O'Regan D, "Learning interpretable anatomical features through deep generative models: Application to cardiac remodeling," In *MICCAI*, 2018.
- [21]. Bishop C, "Improving the generalization properties of radial basis function neural networks," *Neural computation*, vol. 3, no. 4, pp. 579–88, 1991. [PubMed: 31167338]
- [22]. Bozinovski S, Sestakov M, Bozinovska L, "Using EEG alpha rhythm to control a mobile robot," In *Proceedings of the Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pp. 1515–1516, 1988.
- [23]. Bucilua C, Caruana R, Niculescu-Mizil A, "Model compression," In *KDD*, pp. 535–541, 2016.
- [24]. Bychkov D, Linder N, Turkki R, Nordling S, Kovanen PE, Verrill C, Walliander M, Lundin M, Haglund C, Lundin J, "Deep learning based tissue analysis predicts outcome in colorectal cancer," *Scientific reports*, vol. 8, no. 1, pp. 1–1, 2018. [PubMed: 29311619]

- [25]. Caicedo-Torres W and Gutierrez J, "ISeeU: Visually interpretable deep learning for mortality prediction inside the ICU," arXiv preprint, arXiv:1901.08201, 2019.
- [26]. Caruana R, Kangaroo H, Dionisio JD, Sinha U, Johnson D, "Case-based explanation of non-case-based learning methods," In Proceedings of the AMIA Symposium, 1999.
- [27]. Casalicchio G, Molnar C and Bischl B, "Visualizing the feature importance for black box models," In Joint European Conference on Machine Learning and Knowledge Discovery in Databases, pp. 655–670, 2018.
- [28]. Chakraborty S, "Interpretability of deep learning models: a survey of results," IEEE Smart World, Ubiquitous Intelligence & Computing, Advanced & Trusted Computed, Scalable Computing & Communications, Cloud & Big Data Computing, Internet of People and Smart City Innovation, 2017.
- [29]. Che Z, Purushotham S, Khemani R, Liu Y, "Interpretable deep models for ICU outcome prediction," In AMIA Annual Symposium Proceedings, 2016.
- [30]. Chen C, Lin K, Rudin C, Shaposhnik Y, Wang S, Wang T, "An interpretable model with globally consistent explanations for credit risk," arXiv preprint, arXiv:1811.12615, 2018.
- [31]. Chen C, Li O, Barnett A, Su J, Rudin C, "This looks like that: deep learning for interpretable image recognition," arXiv preprint, arXiv:1806.10574, 2018.
- [32]. Chen TQ, Rubanova Y, Bettencourt J & Duvenaud DK, Neural ordinary differential equations. In NeurIPS, pp. 6571–6583, 2018.
- [33]. Chen X, Duan Y, Houthoof R, Schulman J, Sutskever I, and Abbeel P, "Infogan: Interpretable representation learning by information maximizing generative adversarial nets," In NeurIPS, 2016.
- [34]. Chorowski J and Zurada JM, "Learning understandable neural networks with nonnegative weight constraints," IEEE transactions on neural networks and learning systems, vol. 26, no. 1, pp. 62–69, 2014.
- [35]. Chu L, Hu X, Hu J, Wang L and Pei J, "Exact and consistent interpretation for piecewise linear neural networks: A closed form solution," In KDD, pp. 1244–1253, 2018, July.
- [36]. Codella NCF, Lin CC, Halpern A, Hind M, Feris R, and Smith JR, "Collaborative Human-AI (CHAI): Evidence-Based Interpretable Melanoma Classification in Dermoscopic Images." In Understanding and Interpreting Machine Learning in Medical Image Computing Applications: First International Workshops, in MICCAI 2018, Granada, Spain.
- [37]. Cohen N, Sharir O, and Shashua A, "The power of deeper networks for expressing natural functions," In COLT, 2016.
- [38]. Craven MW and Shavlik JW, "Extracting Tree-structured Representations of Trained Networks," In NeurIPS, 1995.
- [39]. Datta A, Sen S, Zick Y, "Algorithmic transparency via quantitative input influence: Theory and experiments with learning systems," In IEEE symposium on security and privacy (SP), pp. 598–617, 2016 May.
- [40]. Devlin J, Chang MW, Lee K, Toutanova K, "Bert: Pre-training of deep bidirectional transformers for language understanding," arXiv preprint, arXiv:1810.04805, 2018.
- [41]. Dosovitskiy A, Brox T, "Inverting visual representations with convolutional networks," In CVPR, pp. 4829–4837, 2016.
- [42]. Doshi-Velez F and Kim B, "Towards a rigorous science of interpretable machine learning," arXiv preprint arXiv:1702.08608, 2017.
- [43]. Du M, Liu N and Hu X, "Techniques for interpretable machine learning," arXiv preprint, arXiv:1808.00033, 2018.
- [44]. Eldan R and Shamir O, "The power of depth for feedforward neural networks," In COLT, 2016.
- [45]. Erhan D, Bengio Y, Courville A, Vincent P, "Visualizing higher-layer features of a deep network," University of Montreal, vol. 1341, no. 3, 2009 Jun 9.
- [46]. Fan F and Wang G, "Learning from Pseudo-Randomness with an Artificial Neural Network—Does God Play Pseudo-Dice?" IEEE Access, vol. 6, pp. 22987–22992, 2018.
- [47]. Fan F and Cong W and Wang G, "A new type of neurons for machine learning," International journal for numerical methods in biomedical engineering, vol. 34, no. 2, e2920, 2018.

- [48]. Fan F and Wang G, “Fuzzy logic interpretation of quadratic networks,” *Neurocomputing*, vol. 374, no. 10–21, 2020.
- [49]. Fan F, Li M, Teng Y and Wang G, “Soft Autoencoder and Its Wavelet Adaptation Interpretation,” *IEEE Transactions on Computational Imaging*, vol. 6, pp. 1245–1257, 2020.
- [50]. Fan L, “Revisit fuzzy neural network: Demystifying batch normalization and ReLU with generalized hamming network,” In *NeurIPS*, pp. 1923–1932, 2017.
- [51]. Fang X, Yan P, “Multi-organ Segmentation over Partially Labeled Datasets with Multi-scale Feature Abstraction,” *arXiv preprint, arXiv:2001.00208*. 2020.
- [52]. Fong RC and Vedaldi A, “Interpretable explanations of black boxes by meaningful perturbation,” In *CVPR*, pp. 3429–3437, 2017.
- [53]. Friedman JH, “Greedy function approximation: a gradient boosting machine,” *Annals of statistics*, pp. 1189–232, 2001.
- [54]. Fu L, “Rule generation from neural networks,” *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 24, no. 8, pp. 1114–24, 1994.
- [55]. Gale W, Oakden-Rayner L, Carneiro G, Bradley AP, and Palmer LJ, “Producing radiologist-quality reports for interpretable artificial intelligence,” *arXiv preprint arXiv:1806.00340*, 2018.
- [56]. Ge R, Huang F, Jin C and Yuan Y, “Escaping from saddle points—online stochastic gradient for tensor decomposition,” *COLT*, 2015.
- [57]. Geis JR, et al. , “Ethics of artificial intelligence in radiology: summary of the joint European and North American multisociety statement,” *Canadian Association of Radiologists Journal*, vol. 70, no. 4, pp. 329–34, 2019. [PubMed: 31585825]
- [58]. Gilpin LH, Bau D, Yuan BZ, Bajwa A, Specter M, Kagal L, “Explaining explanations: An overview of interpretability of machine learning,” In *DSAA*, pp. 80–89, 2018.
- [59]. Goldstein A, Kapelner A, Bleich J, Pitkin E, “Peeking inside the black box: Visualizing statistical learning with plots of individual conditional expectation,” *Journal of Computational and Graphical Statistics*, vol. 24, no. 1, pp. 44–65, 2015.
- [60]. Goodfellow I, Pouget-Abadie J, Mirza M, Xu B, Warde-Farley D, Ozair S, Courville A and Bengio Y, “Generative adversarial nets,” In *NeurIPS*, pp. 2672–2680, 2014.
- [61]. Goodman B and Flaxman S, “European Union regulations on algorithmic decision-making and a ‘right to explanation.’” *AI Mag*, vol. 38, no. 3, pp. 50–57, 2017.
- [62]. Goyal Y, Wu Z, Ernst J, Batra D, Parikh D, and Lee S, “Counterfactual Visual Explanations,” *arXiv preprint, arXiv:1904.07451*, 2019.
- [63]. Guan C, Wang X, Zhang Q, Chen R, He D, Xie X, “Towards a Deep and Unified Understanding of Deep Neural Models in NLP,” In *ICML*, pp. 2454–2463, 2019.
- [64]. Guidotti R, Monreale A, Ruggieri S, Pedreschi D, Turini F, Giannotti F, “Local rule-based explanations of black box decision systems,” *arXiv preprint, arXiv:1805.10820*. 2018.
- [65]. Guidotti R, Monreale A, Ruggieri S, Turini F, Giannotti F & Pedreschi D, “A survey of methods for explaining black box models,” *ACM computing surveys (CSUR)*, vol. 51, no. 5, pp. 93, 2019.
- [66]. Guo H, Kruger U, Wang G, Kalra MK, Yan P, “Knowledge-based Analysis for Mortality Prediction from CT Images,” *arXiv preprint, arXiv:1902.07687*, 2019.
- [67]. Guo Z, Li X, Huang H, Guo N, Li Q, “Deep learning-based image segmentation on multimodal medical imaging,” *IEEE Transactions on Radiation and Plasma Medical Sciences*, vol. 3, no. 2, pp. 162–9, 2019. [PubMed: 34722958]
- [68]. Hatt M, Parmar C, Qi J, EI Naqa I, “Machine (deep) learning methods for image processing and radiomics,” *IEEE Transactions on Radiation and Plasma Medical Sciences*, vol. 3, no. 2, pp. 104–8, 2019.
- [69]. He K, Zhang X, Ren S and Sun Jian, “Deep Residual Learning for Image Recognition”, In *CVPR*, 2016.
- [70]. Hinton GE, Salakhutdinov RR, “Reducing the dimensionality of data with neural networks,” *Science*, vol. 313, no. 5786, pp. 504–7, 2006. [PubMed: 16873662]
- [71]. Hinton GE, Osindero S, The YW, “A fast learning algorithm for deep belief nets,” *Neural computation*, vol. 18, no. 7, pp. 1527–54, 2006. [PubMed: 16764513]

- [72]. Hinton G, Vinyals O, Dean J, “Distilling the knowledge in a neural network,” arXiv preprint, arXiv:1503.02531, 2015.
- [73]. Hochreiter S, Schmidhuber J, “Long short-term memory,” *Neural computation*, vol. no. 8, pp. 1735–80, 1997. [PubMed: 9377276]
- [74]. Hooker G, “Discovering additive structure in black box functions,” In *KDD*, pp. 575–580, 2004.
- [75]. Hu Z, Ma X, Liu Z, Hovy E and Xing E, “Harnessing deep neural networks with logic rules,” arXiv preprint, arXiv:1603.06318., 2016.
- [76]. Huang G, Liu Z, Van Der Maaten L, Weinberger KQ, “Densely connected convolutional networks,” In *CVPR*, pp. 4700–4708, 2017.
- [77]. Huang H, “Mechanisms of dimensionality reduction and decorrelation in deep neural networks,” *Physical Review E*, vol. 98, no. 6, pp. 062313, 2018.
- [78]. Jacot A, Gabriel F, Hongler C, “Neural tangent kernel: Convergence and generalization in neural networks,” In *NeurIPS*, pp. 8571–8580, 2018.
- [79]. Jain S, Wallace BC, “Attention is not explanation,” arXiv preprint arXiv:1902.10186. 2019.
- [80]. Jang JS, “ANFIS: adaptive-network-based fuzzy inference system,” *IEEE transactions on systems, man, and cybernetics*, vol. 23, no. 3, pp. 665–685, 1993.
- [81]. Jiang C, Zhao Y, Chu S, Shen L, and Tu K, “Cold-start and Interpretability: Turning Regular Expressions into Trainable Recurrent Neural Networks”. In *EMNLP*, 2020.
- [82]. Kadanoff LP, “Variational principles and approximate renormalization group calculations,” *Physical Review Letters*, vol. 34, no. 16, pp. 1005, 1975.
- [83]. Kádár A, Chrupała G and Alishahi A, “Representation of linguistic form and function in recurrent neural networks,” *Computational Linguistics*, vol. 43, no. 4, pp. 761–780, 2017.
- [84]. Karpathy A, Fei-Fei L, “Deep visual-semantic alignments for generating image descriptions,” In *CVPR*, pp. 3128–3137, 2015.
- [85]. Karpathy A, Johnson J and Fei-Fei L, “Visualizing and understanding recurrent networks,” arXiv preprint, arXiv:1506.02078, 2015.
- [86]. Kim B, Wattenberg M, Gilmer J, Cai C, Wexler J, Viegas F and Sayres R, “Interpretability beyond feature attribution: Quantitative testing with concept activation vectors (tcav),” arXiv preprint, arXiv:1711.11279, 2017.
- [87]. Kingma DP, Welling M, “Auto-encoding variational bayes,” arXiv preprint arXiv:1312.6114, 2013.
- [88]. Kipf TN, Welling M, “Semi-supervised classification with graph convolutional networks,” arXiv preprint, arXiv:1609.02907, 2016.
- [89]. Koh PW and Liang P, “Understanding black-box predictions via influence functions,” In *ICML*, 2017.
- [90]. Kolodner JL, “An introduction to case-based reasoning,” *Artificial intelligence review*, vol. 1, no. 6, pp. 3–4, 1992.
- [91]. Kolouri S, Yin X, and Rohde GK, “Neural Networks, Hypersurfaces, and Radon Transforms,” arXiv preprint, arXiv:1907.02220, 2019.
- [92]. Krishnan R, Sivakumar G, Bhattacharya P, “Extracting decision trees from trained neural networks,” *Pattern recognition*, vol. 1, no. 32, pp. 12, 2019.
- [93]. Krizhevsky A and Hinton G, “Learning multiple layers of features from tiny images,” 2009.
- [94]. Krotov D and Hopfield JJ, “Unsupervised learning by competing hidden units,” *Proceedings of the National Academy of Sciences*, vol. 116, no. 16, pp. 7723–7731, 2019.
- [95]. Kuo CCCJ, Zhang M, Li S, Duan J and Chen Y, “Interpretable convolutional neural networks via feedforward design,” *Journal of Visual Communication and Image Representation*, vol. 60, pp. 346–359, 2019.
- [96]. Lage I, Ross A, Gershman SJ, Kim B, and Doshi-Velez F, “Human-in-the-loop interpretability prior,” In *NeurIPS*, pp. 10159–10168, 2018.
- [97]. Lakkaraju H, Kamar E, Caruana R and Horvitz E, “Identifying unknown unknowns in the open world: Representations and policies for guided exploration,” In *AAAI*, 2017.
- [98]. Lakkaraju H, Kamar E, Caruana R, Leskovec J, “Interpretable & explorable approximations of black box models,” arXiv preprint, arXiv:1707.01154, 2017.

- [99]. Lark RM, "A comparison of some robust estimators of the variogram for use in soil survey," *European journal of soil science*, vol. 51, no. 1, pp. 137–157, 2000.
- [100]. Lee H, et al. "An explainable deep-learning algorithm for the detection of acute intracranial haemorrhage from small datasets," *Nature Biomedical Engineering*, vol. 3, no. 3, pp. 173, 2019.
- [101]. Lei N, Su K, Cui L, Yau ST, Gu XD, "A geometric view of optimal transportation and generative model," *Computer Aided Geometric Design*, vol. 68, pp. 1–21, 2019.
- [102]. Lei T, Barzilay R and Jaakkola T, "Rationalizing neural predictions," arXiv preprint, arXiv:1606.04155, 2016.
- [103]. Lei Y, et al. "Shape and margin-aware lung nodule classification in low-dose CT images via soft activation mapping," *Medical Image Analysis*, vol. 60, pp. 101628, 2020. [PubMed: 31865281]
- [104]. Li C, et al. , "Deep supervision with intermediate concepts," *IEEE transactions on pattern analysis and machine intelligence*, vol. 41, no. 8, 1828–1843, 2018. [PubMed: 30106706]
- [105]. Li J, Monroe W and Jurafsky D, "Understanding neural networks through representation erasure," arXiv preprint, arXiv:1612.08220, 2016.
- [106]. Li TR, Chamrajnagar A, Fong X, Rizik N, Fu F, "Sentiment-based prediction of alternative cryptocurrency price fluctuations using gradient boosting tree model," *Frontiers in Physics*, vol. 7, no. 98, 2019.
- [107]. Li TY and Yorke JA, "Period three implies chaos," *The American Mathematical Monthly*, vol. 82, no. 10, pp. 985–992, 1975.
- [108]. Li Y, Yosinski J, Clune J, Lipson H, Hopcroft JE, "Convergent Learning: Do different neural networks learn the same representations?" In *ICLR*, 2016.
- [109]. Liang S and Srikant R. Why deep neural networks for function approximation? In *ICLR*, 2017.
- [110]. Lin H and Jegelka S, "Resnet with one-neuron hidden layers is a universal approximator," In *NeurIPS*, 2018.
- [111]. Lindeberg T, "A computational theory of visual receptive fields," *Biological cybernetics*, vol. 107, no. 6, pp. 589–635, 2013. [PubMed: 24197240]
- [112]. Lipton ZC, "The mythos of model interpretability," *Queue*, vol. 16, no. 3, pp. 31–57, 2018.
- [113]. Lipovetsky S and Conklin M, "Analysis of regression in game theory approach," *Applied Stochastic Models in Business and Industry*, vol. 17, no. 4, pp. 319–330, 2001.
- [114]. Lu Y, Zhong A, Li Q, Dong B, "Beyond finite layer neural networks: Bridging deep architectures and numerical differential equations," arXiv preprint, arXiv:1710.10121, 2017.
- [115]. Lundberg SM and Lee SI, "A unified approach to interpreting model predictions," In *NeurIPS*, 2017.
- [116]. Maaten LV, Hinton G, "Visualizing data using t-SNE," *Journal of machine learning research*, pp. 2579–605, 2008.
- [117]. Mahendran A, Vedaldi A, "Understanding deep image representations by inverting them," In *CVPR*, pp. 5188–5196, 2015.
- [118]. Maire F, "On the convergence of validity interval analysis," *IEEE transactions on neural networks*, vol. 11, no. 3, pp. 802–807, 2000. [PubMed: 18249807]
- [119]. Marblestone AH, Wayne G and Kording KP, "Toward an integration of deep learning and neuroscience," *Frontiers in computational neuroscience*, vol. 10, no. 94, 2016.
- [120]. McCulloch WS and Pitts W, "A logical calculus of the ideas immanent in nervous activity," *The bulletin of mathematical biophysics*, vol. 5, no. 4, pp. 115–133, 1943.
- [121]. Mehta P, Schwab DJ, "An exact mapping between the variational renormalization group and deep learning," arXiv preprint, arXiv:1410.3831. 2014.
- [122]. Mei S, Montanari A, "The generalization error of random features regression: Precise asymptotics and double descent curve. arXiv preprint, arXiv:1908.05355. 2019.
- [123]. Melis DA and Jaakkola T, "Towards robust interpretability with self-explaining neural networks," In *NeurIPS*, pp. 7775–7784, 2018.
- [124]. Mhaskar HN and Poggio T, "Deep vs. shallow networks: An approximation theory perspective," *Analysis and Applications*, vol. 14, pp. 829–848, 2016




- [125]. Montavon G, Lapuschkin S, Binder A, Samek W, Müller KR, “Explaining nonlinear classification decisions with deep Taylor decomposition,” *Pattern Recognition*, vol. 65, pp. 211–22, 2017.
- [126]. Nauck D, “A fuzzy perceptron as a generic model for neuro-fuzzy approaches,” *Proc. Fuzzy*, 1994.
- [127]. Neyshabur B, Tomioka R, and Srebro N. Norm-based capacity control in neural networks. In *Conference on Learning Theory*, pp. 1376–1401, 2015.
- [128]. Nguyen A, Dosovitskiy A, Yosinski J, Brox T & Clune J, “Synthesizing the preferred inputs for neurons in neural networks via deep generator networks,” In *NeurIPS*, pp. 3387–3395, 2016.
- [129]. Nguyen A, Clune J, Bengio Y, Dosovitskiy A, Yosinski J, “Plug & play generative networks: Conditional iterative generation of images in latent space,” In *CVPR*, pp. 4467–4477, 2017.
- [130]. Nguyen Q and Hein M, “The loss surface of deep and wide neural networks,” In *ICML*, 2017.
- [131]. Noack A, Ahern I, Dou D, Li B, “Does Interpretability of Neural Networks Imply Adversarial Robustness?” *arXiv preprint*, arXiv:1912.03430, 2019.
- [132]. Oh J, Guo X, Lee H, Lewis RL, Singh S, “Action-conditional video prediction using deep networks in atari games,” In *NeurIPS*, pp. 2863–2871, 2015.
- [133]. Olah C, Mordvintsev A, Schubert L, “Feature visualization,” *Distill*, vol. 2, pp. 11, e7, 2017.
- [134]. Oktay O, et al. , “Attention u-net: Learning where to look for the pancreas,” *arXiv preprint*, arXiv:1804.03999, 2018.
- [135]. Oquab M, Bottou L, Laptev I, Sivic J, “Is object localization for free?-weakly-supervised learning with convolutional neural networks,” In *CVPR*, pp. 685–694, 2015.
- [136]. Park J and Sandberg IW, “Universal approximation using radial-basis-function networks,” *Neural computation*, vol. 3, no. 2, 246–257, 1991. [PubMed: 31167308]
- [137]. Patro B and Namboodiri VP, “Differential attention for visual question answering,” In *CVPR*, pp. 7680–7688, 2018.
- [138]. Pereira S, Meier R, McKinley R, Wiest R, Alves V, Silva CA, Reyes M, “Enhancing interpretability of automatically extracted machine learning features: application to a RBM-Random Forest system on brain lesion segmentation,” *Medical image analysis*, vol. 44, no. 228–44, 2018.
- [139]. Pinheiro PO and Collobert R, “From image-level to pixel-level labeling with convolutional networks,” In *CVPR*, pp. 1713–1721, 2015.
- [140]. Poggio T and Girosi F, “Regularization algorithms for learning that are equivalent to multilayer networks,” *Science*, vol. 247, no. 4945, pp. 978–982, 1990. [PubMed: 17776454]
- [141]. Ribeiro MT, Singh S, Guestrin C, “Why should I trust you?: Explaining the predictions of any classifier,” In *KDD*, pp. 1135–1144, 2016.
- [142]. Ribeiro MT, Singh S, Guestrin C, “Anchors: High-precision model-agnostic explanations,” In *AAAI*, 2018.
- [143]. Robnik-Šikonja M, Kononenko I, “Explaining classifications for individual instances,” *IEEE Transactions on Knowledge and Data Engineering*, vol. 20, no. 5, pp. 589–600, 2008.
- [144]. Rolnick D and Tegmark M., “The power of deeper networks for expressing natural functions,” In *ICLR*, 2018.
- [145]. Ross AS, Hughes MC & Doshi-Velez F, “Right for the right reasons: Training differentiable models by constraining their explanations,” *arXiv preprint*, arXiv:1703.03717, 2017.
- [146]. Rudin C, “Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead,” *Nature Machine Intelligence*, vol. 1, no. 5, pp. 206, 2019.
- [147]. Saad EW, Wunsch II DC, “Neural network explanation using inversion,” *Neural networks*, vol. 20, no. 1, pp. 78–93, 2007. [PubMed: 17029713]
- [148]. Salakhutdinov R, Mnih A, Hinton G, “Restricted Boltzmann machines for collaborative filtering,” In *ICML*, pp. 791–798, 2017
- [149]. Schulz K, Sixt L, Tombari F, Landgraf T, “Restricting the flow: Information bottlenecks for attribution,” In *ICLR*, 2020.
- [150]. Scott T, Ridgeway K, Mozer MC, “Adapted deep embeddings: A synthesis of methods for k-shot inductive transfer learning,” In *NeurIPS*, pp. 76–85, 2018.

- [151]. Selvaraju RR, Cogswell M, Das A, Vedantam R, Parikh D & Batra D, Grad-cam: Visual explanations from deep networks via gradient-based localization. In CVPR, pp. 618–626, 2017.
- [152]. Setiono R and Liu H, “Understanding neural networks via rule extraction,” In IJCAI, vol. 1, pp. 480–485, 2017.
- [153]. Shan H, et al. , “Competitive performance of a modularized deep neural network compared to commercial algorithms for low-dose CT image reconstruction. Nature Machine Intelligence, vol. 1, no. 6, pp. 269, 2019.
- [154]. Shwartz-Ziv R and Tishby N, “Opening the black box of deep neural networks via information,” arXiv preprint, arXiv:1703.00810, 2017.
- [155]. Shen S, Han SX, Aberle DR, Bui AA and Hsu W, “An interpretable deep hierarchical semantic convolutional neural network for lung nodule malignancy classification,” Expert Systems with Applications, vol. 128, pp. 84–95, 2019. [PubMed: 31296975]
- [156]. Shrikumar A, Greenside P, Shcherbina A and Kundaje A, “Not just a black box: Interpretable deep learning by propagating activation differences,” In ICML, 2016.
- [157]. Simonyan K, Vedaldi A and Zisserman A, “Deep inside convolutional networks: Visualising image classification models and saliency maps,” arXiv preprint, arXiv:1312.6034, 2013.
- [158]. Simonyan K, Zisserman A, “Very deep convolutional networks for large-scale image recognition,” In CVPR, 2015.
- [159]. Singh C, Murdoch WJ, Yu B, “Hierarchical interpretations for neural network predictions,” In ICLR, 2019.
- [160]. Singla S, Wallace E, Feng S and Feizi S, “Understanding Impacts of High-Order Loss Approximations and Features in Deep Learning Interpretation,” arXiv preprint, arXiv:1902.00407, 2019.
- [161]. Smilkov D, Thorat N, Kim B, Viégas F and Wattenberg M, “Smoothgrad: removing noise by adding noise,” arXiv preprint, arXiv:1706.03825, 2017.
- [162]. Snyder SH, “Adenosine as a neuromodulator,” Annual review of neuroscience, vol. 8, no. 1, pp. 103–24, 1985.
- [163]. Soltanolkotabi M, Javanmard A and Lee JD, “Theoretical insights into the optimization landscape of over-parameterized shallow neural networks,” IEEE Transactions on Information Theory, vol. 65, no. 2, pp. 742–769, 2018.
- [164]. Springenberg JT, Dosovitskiy A, Brox T and Riedmiller M, “Striving for simplicity: The all convolutional net,” arXiv preprint, arXiv:1412.6806, 2014.
- [165]. Stone A, Wang H, Stark M, Liu Y, Phoenix DS and George D, “Teaching compositionality to cnns,” In CVPR, pp. 5058–5067, 2017.
- [166]. Sturm I, Lapuschkin S, Samek W and Müller KR, “Interpretable deep neural networks for single-trial EEG classification,” Journal of neuroscience methods, vol. 274, pp. 141–145, 2016. [PubMed: 27746229]
- [167]. Subramanian A, Pruthi D, Jhamtani JH, Berg-Kirkpatrick T and Hovy E, “Spine: Sparse interpretable neural embeddings,” In AAAI, 2018.
- [168]. Sundararajan M, Taly A and Yan Q, “Axiomatic attribution for deep networks,” In ICML, pp. 3319–3328, 2017.
- [169]. Szegedy C, Zaremba W, Sutskever I, Bruna J, Erhan D, Goodfellow I, Fergus R, “Intriguing properties of neural networks,” arXiv preprint, arXiv:1312.6199, 2013.
- [170]. Szymanski L and McCane B. Deep networks are effective encoders of periodicity. IEEE transactions on neural networks and learning systems, 25:1816–1827, 2014. [PubMed: 25291735]
- [171]. Tan S, Caruana R, Hooker G, Koch P, Gordo A, “Learning Global Additive Explanations for Neural Nets Using Model Distillation,” arXiv preprint, arXiv:1801.08640, 2018.
- [172]. Takagi T, Sugeno M, “Fuzzy identification of systems and its applications to modeling and control,” IEEE transactions on systems, man, and cybernetics, pp. 116–32, 1985.
- [173]. Thrun S, “Extracting rules from artificial neural networks with distributed representations,” In NeurIPS, pp. 505–512, 1995.

- [174]. Torres-Velázquez M, Chen WJ, Li X, McMillan AB, “Application and Construction of Deep Learning Networks in Medical Imaging,” *IEEE Transactions on Radiation and Plasma Medical Sciences*, 2020 Oct 13.
- [175]. Van der Maas HL, Verschure PF & Molenaar PC, “A note on chaotic behavior in simple neural networks,” *Neural Networks*, vol. 3, no. 1, pp. 119–122, 1990.
- [176]. Van Molle P, De Strooper M, Verbelen T, Vankeirsbilck B, Simoons P, Dhoedt B, “Visualizing convolutional neural networks to improve decision support for skin lesion classification,” In *Understanding and Interpreting Machine Learning in Medical Image Computing Applications*, pp. 115–123, 2018.
- [177]. Vaughan J, Sudjianto A, Brahimi E, Chen J, and Nair VN, “Explainable neural networks based on additive index models,” *arXiv preprint, arXiv:1806.01933*, 2018.
- [178]. Veit A, Wilber MJ, and Belongie S. Residual networks behave like ensembles of relatively shallow networks. In *NeurIPS*, 2016.
- [179]. Vinyals O, Toshev A, Bengio S and Erhan D, “Show and tell: A neural image caption generator,” In *CVPR*, pp. 3156–3164, 2015.
- [180]. Wachter S, Mittelstadt B and Russell C, “Counterfactual Explanations without Opening the Black Box: Automated Decisions and the GDPR,” *Harv. JL & Tech*, vol. 31, no. 841, 2017.
- [181]. Wallace E, Feng S, Boyd-Graber J, “Interpreting Neural Networks with Nearest Neighbors,” *arXiv preprint, arXiv:1809.02847*, 2018.
- [182]. Wang G, “A perspective on deep imaging,” *IEEE Access*, vol. 4, pp. 8914–24, 2016.
- [183]. Wang T, “Gaining Free or Low-Cost Interpretability with Interpretable Partial Substitute,” In *ICML*, pp. 6505–6514, 2019.
- [184]. Wang Y, Su H, Zhang B and Hu X, “Interpret neural networks by identifying critical data routing paths,” In *CVPR*, 2018.
- [185]. Worrall DE, Garbin SJ, Turmukhambetov D and Brostow GJ, “Interpretable transformations with encoder-decoder networks,” In *ICCV*, pp. 5726–5735, 2017.
- [186]. Wu M, Hughes MC, Parbhoo S, Zazzi M, Roth V and Doshi-Velez F, “Beyond sparsity: Tree regularization of deep models for interpretability,” In *AAAI*, 2018.
- [187]. Wu T, Sun W, Li X, Song X & Li B, “Towards Interpretable R-CNN by Unfolding Latent Structures,” *arXiv preprint, arXiv:1711.05226*. 2017.
- [188]. Wu W, Hu D, Wang S, Yu H, Vardhanabhuti V, Wang G, “Stabilizing Deep Tomographic Reconstruction Networks,” *arXiv preprint, arXiv:2008.01846*. 2020.
- [189]. Xie Q, Ma X, Dai Z and Hovy E, “An interpretable knowledge transfer model for knowledge base completion,” *arXiv preprint, arXiv:1704.05908*, 2017.
- [190]. Xu K, Ba J, Kiros R, Cho K, Courville A, Salakhudinov R, Zemel R, Bengio Y, “Show, attend and tell: Neural image caption generation with visual attention,” In *ICML*, pp. 2048–2057, 2015.
- [191]. Xu K, Liu S, Zhang G, et al. , “Interpreting adversarial examples by activation promotion and suppression,” *arXiv preprint, arXiv:1904.02057*, 2019.
- [192]. Yang C, Rangarajan A and Ranka S, Global model interpretation via recursive partitioning. In *IEEE 20th International Conference on High Performance Computing and Communications; IEEE 16th International Conference on Smart City; IEEE 4th International Conference on Data Science and Systems (HPCC/SmartCity/DSS)*, pp. 1563–1570, 2018, June.
- [193]. Ye JC, Han Y, and Cha E, “Deep convolutional framelets: A general deep learning framework for inverse problems,” *SIAM Journal on Imaging Sciences*, vol. 11, no. 2, pp. 991–1048, 2018.
- [194]. Ying Z, Bourgeois D, You J, Zitnik M, Leskovec J, “Gnnexplainer: Generating explanations for graph neural networks,” In *Advances in neural information processing systems*, pp. 9244–9255, 2019.
- [195]. You S, Ding D, Canini K, Pfeifer J and Gupta M, “Deep lattice networks and partial monotonic functions,” In *NeurIPS*, pp. 2981–2989, 2017.
- [196]. You J, Leskovec J, He K, and Xie S, “Graph Structure of Neural Networks,” In *ICML*, 2020.
- [197]. Yosinski J, Clune J, Nguyen A, Fuchs T, Lipson H, “Understanding neural networks through deep visualization,” *arXiv preprint, arXiv:1506.06579*. 2015.

- [198]. Yu S and Principe JC, "Understanding autoencoders with information theoretic concepts," *Neural Networks*, vol. 117, pp. 104–23, 2019. [PubMed: 31132606]
- [199]. Zadeh LA, "Fuzzy logic," *Computer*, vol. 21, no. 4, pp. 83–93, 1988.
- [200]. Zech JR, Badgeley MA, Liu M, Costa AB, Titano JJ, & Oermann EK, "Variable generalization performance of a deep learning model to detect pneumonia in chest radiographs: a cross-sectional study," *PLoS medicine*, vol. 15, no. 11, e1002683, 2018. [PubMed: 30399157]
- [201]. Zeiler MD and Fergus R, "Visualizing and understanding convolutional networks," In *ECCV*, pp. 818–833, 2014.
- [202]. Zhang Q and Zhu SC, "Visual interpretability for deep learning: a survey," *Frontiers of Information Technology & Electronic Engineering*, vol. 19, no. 1, pp. 27–39, 2018.
- [203]. Zhang Q, Cao R, Shi F, Wu YN and Zhu SC, "Interpreting cnn knowledge via an explanatory graph," In *AAAI*, 2018.
- [204]. Zhang Q, Wang W and Zhu SC, "Examining cnn representations with respect to dataset bias," In *AAAI*, 2018.
- [205]. Zhang Q, Wu YN and Zhu SC, "Interpretable convolutional neural networks," In *CVPR*, pp. 8827–8836, 2018.
- [206]. Zhang P, Wang J, Farhadi A, Hebert M & Parikh D, "Predicting failures of vision systems," In *CVPR*, pp. 3566–3573, 2014.
- [207]. Zhang Y, Song K, Sun Y, Tan S, Udell M, "Why Should You Trust My Explanation?" *Understanding Uncertainty in LIME Explanations*, arXiv preprint, arXiv:1904.12991. 2019.
- [208]. Zhang Z, Xie Y, Xing F, McGough M, Yang L, "Mdnet: A semantically and visually interpretable medical image diagnosis network," In *CVPR*, pp. 6428–6436, 2017.
- [209]. Zhang Z, Chen P, McGough M, Xing F, Wang C, Bui M, Xie Y, Sapkota M, Cui L, Dhillon J, Ahmad N, "Pathologist-level interpretable whole-slide cancer diagnosis with deep learning," *Nature Machine Intelligence*, vol. 1, no. 5, pp. 236–45, 2019.
- [210]. Zhou B, Khosla A, Lapedriza A, Oliva A and Torralba A, "Learning deep features for discriminative localization," In *CVPR*, pp. 2921–2929, 2016.
- [211]. Zhou B, Khosla A, Lapedriza A, Oliva A and Torralba A, "Object detectors emerge in deep scene cnns," arXiv preprint, arXiv:1412.6856. 2014.
- [212]. Zintgraf LM, Cohen TS, Adel T and Welling M, "Visualizing deep neural network decisions: Prediction difference analysis," arXiv preprint, arXiv:1702.04595, 2017.



**Fig. 1.** Exponential growth of the number of articles on interpretability.

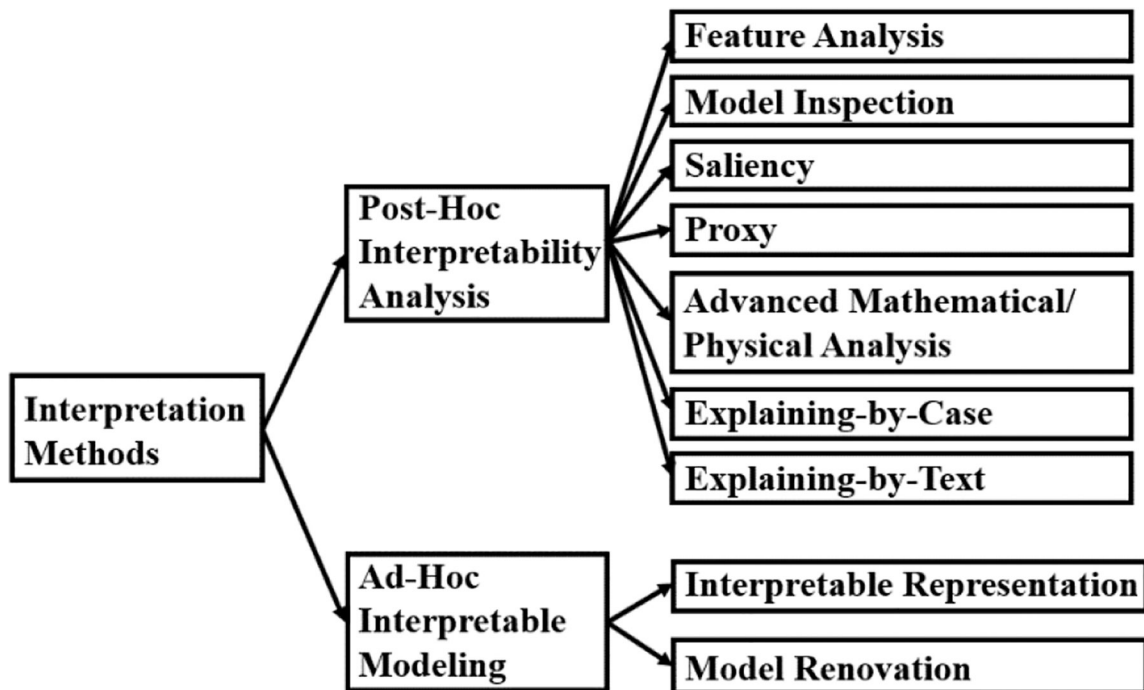

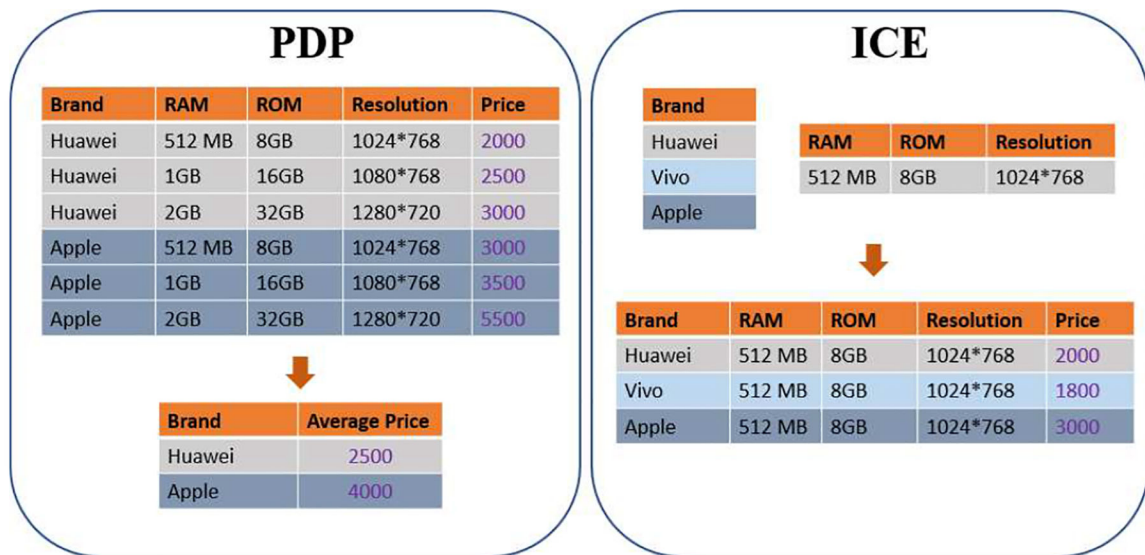


Fig. 2. Taxonomy used for this interpretability review.

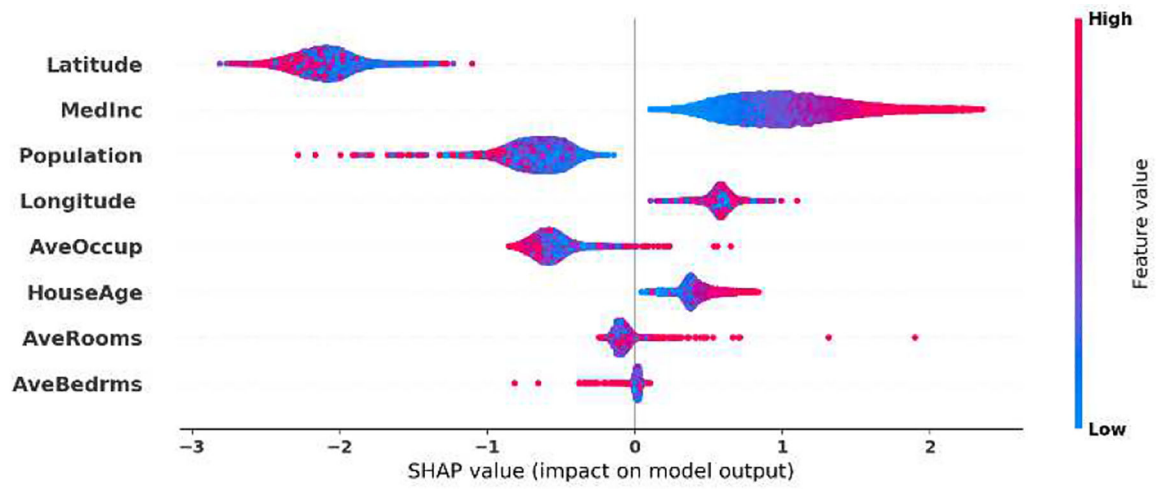


**Fig. 3.** Based on the influence function, two harmful images that have the same label as the test image are identified.

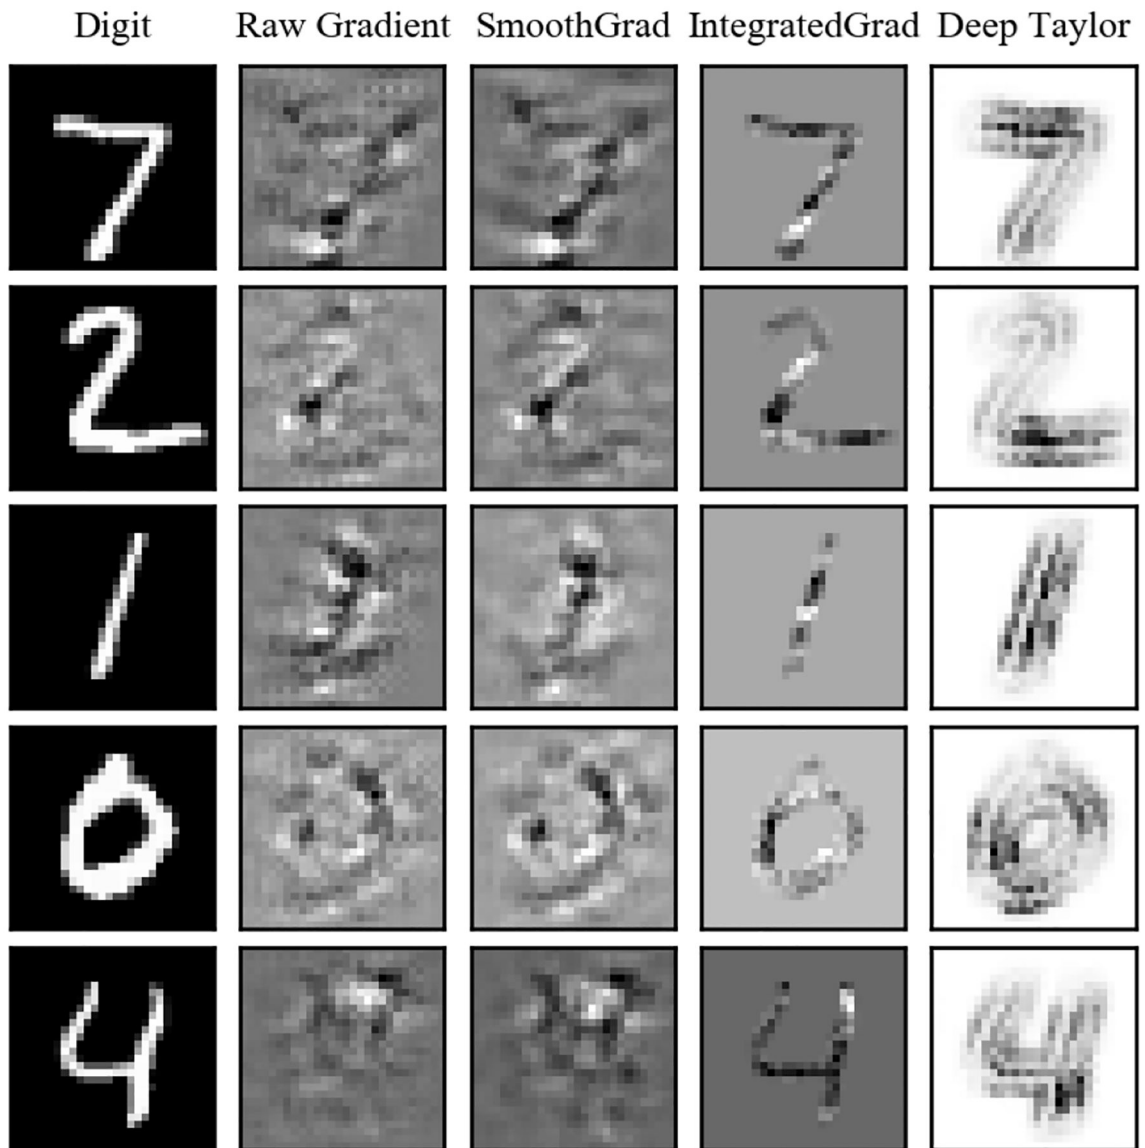


**Fig. 4.** Toy examples illustrating the definitions of PDP and ICE, respectively. On the left, to measure the impact of the brand on the price with the PDP method, we fix the brand and compute the average of prices as other factors change, obtaining that the PDP of “Huawei” is 2500 and the PDP of “Apple” is 4000. On the right, ICE scores regarding brands “Huawei”, “Vivo” and “Apple” are computed by varying brands and fixing other factors.

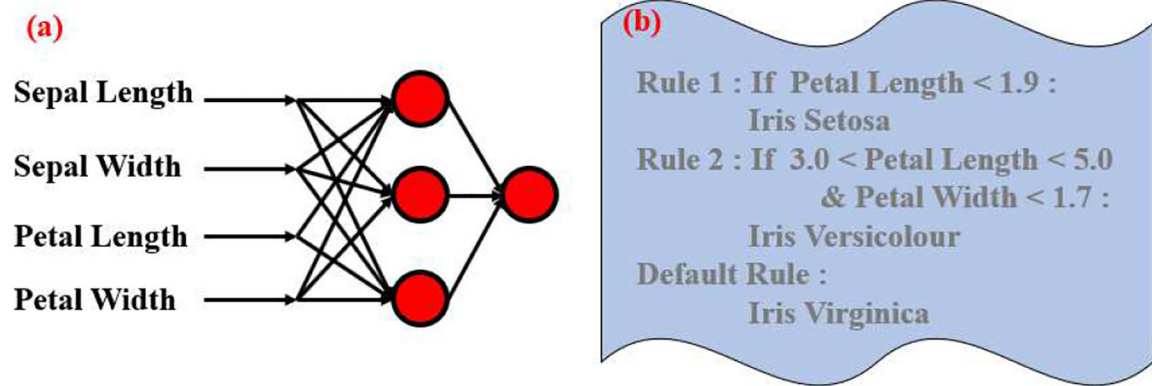





**Fig. 5.** Positive Shapley value indicates a positive impact on the model output, and vice versa. Shapley value analysis shows that the model is biased because the house age has the positive Shapley value on the house price, which goes against with our real experience.




**Fig. 6.** Interpreting a LeNet-5-like network by raw gradient, SmoothGrad, Integrated Gradient, and Deep Taylor methods, respectively. It is seen that Integrated Gradient and Deep Taylor methods have sharper and less noisy saliency map.




**Fig. 7.** Rule extraction process as proposed by R. Setiono and H. Liu [152]. (a) A one-hidden-layer network with three hidden neurons is constructed to classify the Iris dataset. (b) Rules are extracted via discretizing activation values of hidden units and clustering of inputs, where Petal length and Petal width are dominating attributes for classification of Iris samples. The extracted rules have the same classification performance as that of the original neural network.




**Fig. 8.** Knowledge distillation is to construct an interpretable proxy by the soft labels from the original complex models.




**Fig. 9.** A breast cancer classification task model dissected by LIME. In this case, the sample is classified as benign where worst concave point, mean concave point and so on are contributing forces, while the worst perimeter is the contributing force to drive the model to predict “malignant”.



**Fig. 10.**  
ODE-Net optimizes the start point and the dynamics to fit the spiral shape.

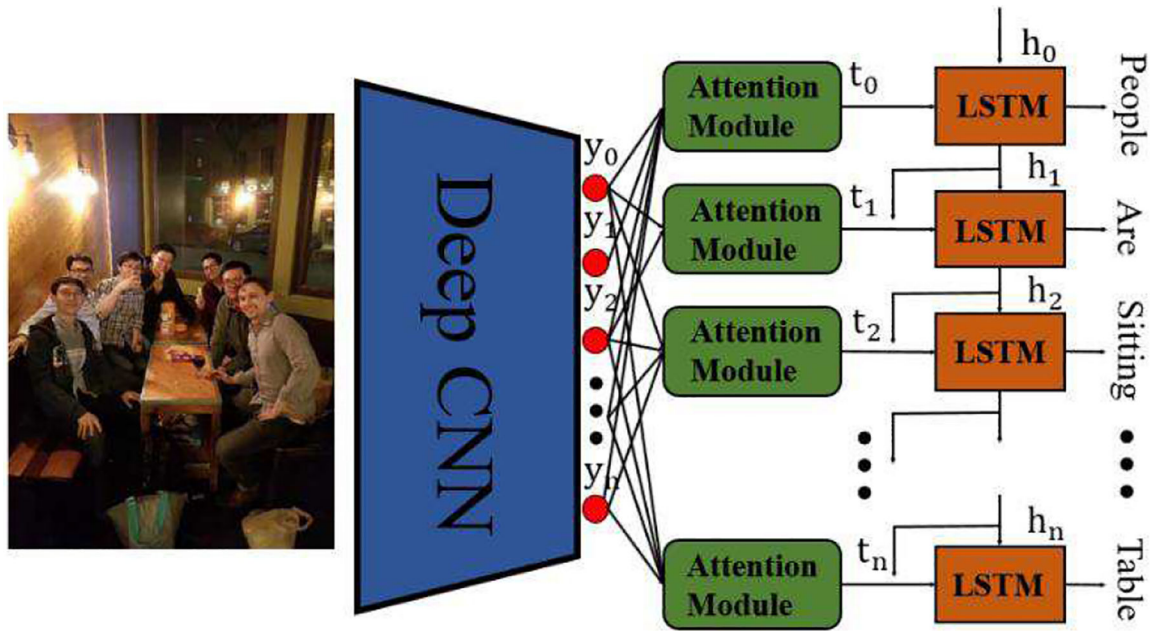


**Fig. 11.**  
An application of information bottleneck theory to compare mutual information between symmetric layers in an autoencoder.




**Fig. 12.** Explaining-by-case presents the nearest neighbors in response to a query.







**Fig. 13.** Image captioning with attention modules provides an explanation to the features mined by a deep convolutional network.




**Fig. 14.**  
In an *InfoGAN*, two latent codes control the localized parts and rotation parts respectively.



**Fig. 15.** Soft-autoencoder with soft-thresholding functions as activation functions in the encoding layers and the linear function as activations in the decoding layers, thereby admitting a direct correspondence to the wavelet adaptation system.



**Fig. 16.** Visualization of feature maps of different arms in PIPO-FAN, where low-scale sub-networks produce local structural details and high-scale sub-networks target global morphological information.



**Fig. 17.** Visualization of weights of a network learned by a bio-plausible algorithm, where prototypes of training image are captured [94].