

# Transparency of deep neural networks for medical image analysis

Citation for published version (APA):

Salahuddin, Z., Woodruff, H. C., Chatterjee, A., & Lambin, P. (2022). Transparency of deep neural networks for medical image analysis: A review of interpretability methods. *Computers in Biology and Medicine*, 140, Article 105111. <https://doi.org/10.1016/j.combiomed.2021.105111>

## Document status and date:

Published: 01/01/2022

## DOI:

[10.1016/j.combiomed.2021.105111](https://doi.org/10.1016/j.combiomed.2021.105111)

## Document Version:

Publisher's PDF, also known as Version of record

## Document license:

CC BY

## Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

## General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

[www.umlib.nl/taverne-license](http://www.umlib.nl/taverne-license)

## Take down policy

If you believe that this document breaches copyright please contact us at:

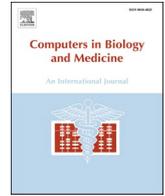
[repository@maastrichtuniversity.nl](mailto:repository@maastrichtuniversity.nl)

providing details and we will investigate your claim.



Contents lists available at ScienceDirect

# Computers in Biology and Medicine

journal homepage: [www.elsevier.com/locate/combiomed](http://www.elsevier.com/locate/combiomed)

## Transparency of deep neural networks for medical image analysis: A review of interpretability methods

Zohaib Salahuddin<sup>a,\*</sup>, Henry C. Woodruff<sup>a,b</sup>, Avishek Chatterjee<sup>a</sup>, Philippe Lambin<sup>a,b</sup><sup>a</sup> The D-Lab, Department of Precision Medicine, GROW – School for Oncology and Developmental Biology, Maastricht University, Maastricht, the Netherlands<sup>b</sup> Department of Radiology and Nuclear Medicine, GROW – School for Oncology and Developmental Biology, Maastricht University Medical Centre+, Maastricht, the Netherlands

### ARTICLE INFO

#### Keywords:

Explainable artificial intelligence  
 Medical imaging  
 Explainability  
 Interpretability  
 Deep neural networks

### ABSTRACT

Artificial Intelligence (AI) has emerged as a useful aid in numerous clinical applications for diagnosis and treatment decisions. Deep neural networks have shown the same or better performance than clinicians in many tasks owing to the rapid increase in the available data and computational power. In order to conform to the principles of trustworthy AI, it is essential that the AI system be transparent, robust, fair, and ensure accountability. Current deep neural solutions are referred to as black-boxes due to a lack of understanding of the specifics concerning the decision-making process. Therefore, there is a need to ensure the interpretability of deep neural networks before they can be incorporated into the routine clinical workflow. In this narrative review, we utilized systematic keyword searches and domain expertise to identify nine different types of interpretability methods that have been used for understanding deep learning models for medical image analysis applications based on the type of generated explanations and technical similarities. Furthermore, we report the progress made towards evaluating the explanations produced by various interpretability methods. Finally, we discuss limitations, provide guidelines for using interpretability methods and future directions concerning the interpretability of deep neural networks for medical imaging analysis.

### 1. Introduction

Medical imaging plays a key role in modern medicine as it allows for the non-invasive visualization of internal structures and metabolic processes of the human body in detail. This aids in disease diagnostics, treatment planning, and treatment follow-up by adding potentially informative data in the form of patient-specific disease characteristics [2,74]. The amount of healthcare imaging data is rapidly increasing due to advances in hardware, the increase in population, decrease in cost, and the awareness of the utility of the imaging modalities [133]. This adds to the increasing difficulty for radiologists and clinicians to cope with the mounting burden of analyzing the large amounts of available data from disparate data sources, and studies have highlighted sometimes considerable inter-observer variability when performing various clinical imaging tasks [112]. It follows that there is an evolving need for tools that can aid in diagnosis and decision-making.

Recent advances in artificial intelligence (AI) have started permeating into healthcare, among those the so-called deep learning (DL) methods, which consist of non-linear modules that can learn multiple

levels of representations automatically from high dimensional data without any need of explicit feature engineering by humans [75,120]. Deep learning methods, in particular convolutional neural networks (CNN), first rose to prominence in 2012 when they demonstrated state-of-the-art performance on the ImageNet Large-Scale Visual Recognition Challenge (ILSVRC) [70], and have since become the *de facto* standard for solving computer vision problems. Deep learning has also demonstrated state-of-the-art performance on many medical imaging challenges related to classification [3,14], segmentation [17,158,173] and other tasks [113].

Despite their promising performance, the adoption of deep learning networks in healthcare has been slow. A contributing factor to this is that more insight into the inner workings of AI tools is needed before translation into the clinic is widely successful [87]. DL models are referred to as black-boxes due to their opaque nature as the interpretation of the inner states of the model is not as straightforward as with standard object-oriented code or decision trees [32]. Furthermore, the European Union's General Data Protection Regulation (GDPR) law requires an explanation of the decision-making process of the algorithm to

\* Corresponding author.

E-mail address: [z.salahuddin@maastrichtuniversity.nl](mailto:z.salahuddin@maastrichtuniversity.nl) (Z. Salahuddin).

<https://doi.org/10.1016/j.combiomed.2021.105111>

Received 31 October 2021; Received in revised form 1 December 2021; Accepted 2 December 2021

Available online 4 December 2021

0010-4825/© 2021 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

make it transparent before it can be utilized for patient-care [140].

The primary purpose of AI tools for medical imaging is to aid clinicians in their decision-making by combining multiple factors into a model that returns an actionable output. Without any explanation of this output, the utility of the model is limited as it does not unveil the reasoning process, limitations, and biases. Interpretability of DL systems cannot only unravel any faulty processes within the algorithms, but also enables the discovery of other important information in the imaging data that otherwise might go unnoticed. Unraveling the black-box nature of the DL systems is not only a legal and ethical requirement but is also essential for fostering clinical trust and for troubleshooting systems. Moreover, interpretability methods can also reveal new imaging biomarkers to understand the specifics of the DL model.

Explainable artificial intelligence (XAI) refers to AI solutions that can provide some details about their functioning in a way that is understandable to the end-users [9]. As the field is still in its infancy, there is no consensus on the definition of the terms interpretability and explainability in this context, and recently various definitions have been proposed for these terms [9,88,110]. For this work, we do not differentiate between interpretability and explainability as in Ref. [18]. We define the interpretability of deep neural networks simply as an attempt to explain the decision-making process of the model in a way that is understandable for the end-users. It refers to any technique that attempts to answer the question “Why is the model making this prediction?” for the medical image analysis tasks.

This review is structured as follows. In section 2, we review and discuss the interpretability methods for understanding DL models trained to solve medical image analysis problems. Section 3 discusses how the explanations produced by various methods can be evaluated to ensure that they are robust, consistent, and useful to clinicians. In section 4, we discuss the shortcomings of interpretability methods, provide guidelines for using interpretability methods, and discuss steps for the future. The objective of the review paper is to discuss the advantages, limitations, and evaluation of interpretability methods for DL models and this will ultimately help in choosing the correct interpretability methods.

This review paper aims to answer the following research questions:

- How interpretability methods can help in the incorporation of deep neural networks in the clinical workflow?
- What are the interpretability methods available for understanding deep neural networks for medical image analysis tasks?
- How can various interpretability methods be grouped into categories?
- What are the advantages, limitations, and challenges of different interpretability methods?
- What are different evaluation methods available to ensure that the explanations are trustworthy?

The following keywords were used for literature search on Google Scholar and PubMed to find answers to the above-mentioned research questions: “interpretability”, “explainability”, “medical imaging”, “Explainable artificial intelligence”, “deep learning”, “evaluation”, “sanity checks”, “clinical validation”, “imaging biomarkers”. The articles were chosen based on their novelty, clinical application, and relevance to clinical workflow.

## 2. Interpretability methods

The purpose of AI solutions is to aid clinicians in performing their work more efficiently and accurately, and not to replace them. This partnership requires trust on the side of the clinical experts, and one path to trust is understanding [87]. We can incorporate interpretability during the design process of the deep neural network. Post-hoc interpretability methods provide explanations for the predictions after the DL model has been trained. Interpretability methods can offer local or global explanations. Local explanations identify the attributes and features of a particular image that the DL model considers important for prediction. On the other hand, global explanations aim at identifying the common characteristics that the DL model considers when associating images with that particular class. After the literature review using systematic keyword searches, we have identified nine different types of interpretability methods for medical image analysis tasks based on the type of generated explanations and the technical similarities. Table 1 lists the different interpretability methods, their limitations, and the application of these methods for interpreting DL models. Fig. 1 presents

**Table 1**

The table summarizes different interpretability methods, their limitations and the application of these methods for explaining DL models for medical image analysis tasks.

No.	Interpretability Methods	Description	Limitations	References
1.	Concept Learning Models	High-level clinical concepts are first predicted and the final classification is made using these concepts.	Additional annotation cost, learned concepts may encode information beyond the intended clinical concepts due to information leakage.	[69,111,125,125]
2.	Case-Based Models	Class discriminative prototypes are learned and the final classification is performed by comparing features extracted from input images with the prototypes.	Susceptibility to corruption by noise and compression artefacts, difficult to train.	[20,66,80,96,129,130]
3.	Counterfactual Explanation	Input images are perturbed in a realistic manner to generate the opposite prediction.	Possibility of unrealistic perturbations to the input images, the resolution of the generated counterfactual images is limited.	[10,12,24,78,121,122,124,128,131,138]
4.	Concept Attribution	Global explanations to quantify the influence of high-level image concepts/radiomics features on the model predictions.	Difficult to annotate high-level clinical concepts, radiomics features used for interpretability may not be reproducible.	[23,23,25,40,49,58,65,159]
5.	Language Description	Textual justifications are provided along with the predictions.	Structured diagnostic reports require more annotation efforts, duplication of training sentences during testing.	[22,76,117,149,165,166]
6.	Latent Space Interpretation	The latent space is used to uncover the salient factors of variation learned in the data with respect to the clinical knowledge. Visualization of high-dimensional latent space in two dimensions to identify similarities and outliers.	Loss of information when the high-dimensional feature space is projected to two dimensions. The similarity in latent space does not always translate to the similarity in terms of human-interpretable features.	[13,21,23,30,38,53,92,106,155,157]
7.	Attribution Map	Post-hoc explanations are provided by highlighting the regions of the input image that the model considers important.	No information is provided on how the relevant regions contribute to the prediction, multiple classes can have the same regions highlighted.	[15,16,44,57,64,90,93,104,118,127,152]
8.	Anatomical Prior	Task-specific structural information is incorporated in the design process of the network.	Specialized clinical knowledge may be required, anatomical prior cannot be utilized for all problems.	[41,105,136]
9.	Internal Network Representation	Visualization of features learned by different filters in a CNN.	The structures and patterns that different filters learn to identify are hard to interpret in medical images.	[11,99]

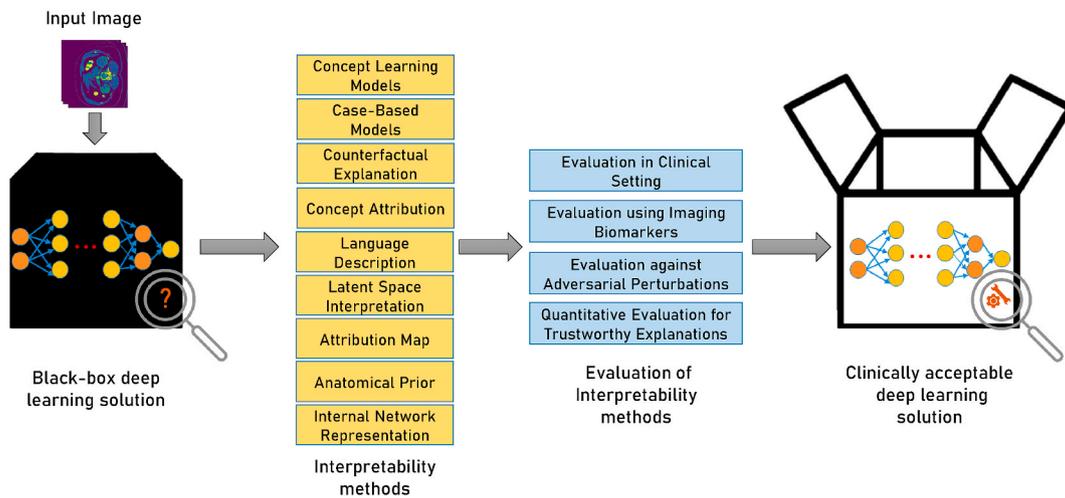


Fig. 1. Overview of Interpretability methods for DL solutions in medical image analysis. The black-box DL solution can be made more desirable for clinical use by incorporating interpretability during the design or execution phase. The evaluation of explanations is important to ensure the robustness of the interpretability methods.

an overview of interpretability methods for DL solutions in medical image analysis.

### 2.1. Concept learning models

State-of-the-art DL models are trained to infer the label  $y$  directly from the input image  $x$ . It is usually not possible for radiologists to understand the reason behind the prediction of the DL models using the same high-level concepts  $c$  used to arrive at a diagnosis. Hence, it is advantageous to explain the outputs of the model in terms of human-interpretable concepts. This problem can be solved by first predicting these high-level concepts (such as semantic features) from the input image and then using these concepts to predict the label [69]. These models require concepts generated by experts as input during training time along with the image and label. Conceptual alignment deep neural networks (CADNNs) utilize some hidden neurons to learn human-interpretable concepts while other neurons are trained freely. CADNNs achieve performance comparable to deep neural networks trained without interpretability constraints [27]. Dai and Wang [26] used a conceptual alignment deep autoencoder to analyze tongue images that represent different body constituent types based on traditional Chinese medicine principles. Koh et al. [69] introduced Concept Bottleneck Models for osteoarthritis grading and used 10 clinical concepts such as joint space narrowing, bone spurs, calcification, etc. The clinicians can intervene at test time to change the predicted value of the clinical concept to observe the effect on the final prediction. This intervention from the clinicians resulted in a performance improvement. This added interpretability comes at a cost of annotation of these clinical concepts, which requires time. Fig. 2 shows how concept learning models first predict clinical concepts and the final prediction is based on either only clinical concepts or on a combination of clinical concepts and other deep features. Margeloiu et al. [95] utilized post-hoc attribution methods to show that the concepts learned in concept bottleneck models do not relate with the input image in a meaningful way. Mahinpei et al. [94] showed that concept bottleneck models could be misleading as the learned encoding contains information in addition to the concept representation.

Capsule networks encode information (e.g. pose, scale, etc.) about each feature using vectorized representation in contrast to scalar features maps used in CNN [111]. LaLonde et al. [71] proposed a multi-task capsule architecture called X-Caps for encoding high-level visual attributes within the capsule vectors to perform explainable diagnosis for lung nodule malignancy prediction. Malignancy prediction was based

entirely on concepts such as sphericity, margin, texture, etc. The predicted visual attributes are regularized by a segmentation reconstruction error. The proposed 2D explainable capsule network outperforms Hierarchical Semantic Convolutional Neural Network by Shen et al. [125] in terms of visual attribute modeling and malignancy score prediction. The performance of X-Caps approaches that of 3D non-explainable CNNs for malignancy prediction.

Low-level features extracted from CNN layers can be combined with predicted clinical concepts for diagnosis. Hierarchical Semantic Convolutional Neural Network (HSCNN) for lung nodule malignancy classification consists of three modules [125]. The first module extracts generalized low-level features that are fed to the second module. The second module classifies the presence or absence of five nodule semantic characteristics that reflect diagnostic features relevant to radiologists (e.g. texture, margin). The third module predicts lung nodule malignancy based on the low-level features from the first module and high-level visual attributes from the second module. The authors showed that HSCNN achieves better results than some non-explainable 3D CNNs.

### 2.2. Case-based models

Case-Based Models make predictions by comparing the features extracted from an input image against class discriminative prototypes. Prototype classification is a type of case-based reasoning that is inherently interpretable because the final predictions are made by taking a weighted sum of similarity scores between features extracted from input and prototypes [80]. Chen et al. [20] proposed a Prototypical Part Network (ProtoPNet) that consists of a convolutional layer, followed by a prototype layer and then a fully connected layer. The convolutional layer consists of a trimmed standard CNN pipeline that acts as a feature extractor. The prototype layer takes patches from the convolution layer as input and learns class discriminative prototypes during training. A similarity score is computed after comparison against each prototype. A fixed-size feature map is used for comparison with prototypes. The fully connected layer then makes predictions based on these similarity scores. The convolutional layer and the prototype layer are trained first and the loss function comprises misclassification loss, cluster cost, and separation cost. In the second step, the fully connected layer is trained. The interpretability of ProtoPNet does not come at a cost of performance when compared to black-box DL models. Mohammadjafari et al. [96] utilized ProtoPNet with DenseNet121 architecture for Alzheimer's disease classification. Barnett et al. [8] used a modified version of ProtoPNet to utilize fine-grained expert annotations for mass margin

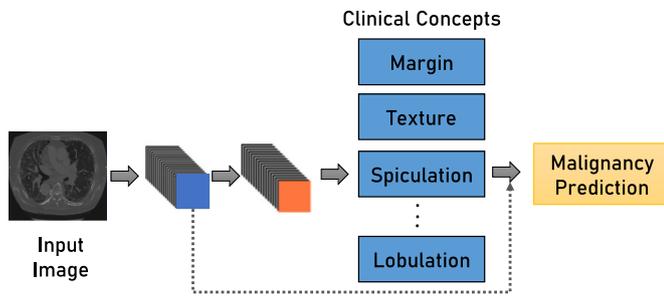


Fig. 2. Concept Learning Models first predict clinical concepts and the final prediction is made using these human interpretable concepts. The final prediction is based either only on the clinical concepts or a combination of clinical concepts and other deep features (dashed line).

classification and malignancy prediction. The prototype networks are comparatively difficult to train than regular deep neural networks because of the scarcity of available techniques for their training [110].

ProtoPNet uses a fixed size feature map from an input image for comparison with prototypes. Hence, there is a possibility that the patch may not be able to capture class discriminative features. XProtoNet extracts features within a dynamic area and is not constrained to a limited size for comparison with prototypes [66]. XProtoNet achieved state-of-the-art performance for the classification of 14 diseases on the public chest X-ray dataset. Generalized Prototypical Part Network (Gen-ProtoPNet) puts no constraint on the dimensions of the prototype [130]. Gen-ProtoPNet was used for the classification of Normal, COVID, and Pneumonia classes in chest X-rays. Negative-Positive Prototypical Part Network (NP-ProtoPNet) is an extension of ProtoPNet that takes the association between similarity scores and incorrect class into consideration [129]. NP-ProtoPNet was utilized for the classification of COVID and Pneumonia in chest X-ray. Similarity scores in ProtoPNet can be corrupted by noise and JPEG compression artefacts [54]. The similarity in latent space does not always translate to the similarity in terms of human-interpretable features.

### 2.3. Counterfactual explanation

A counterfactual explanation is an image that is produced by applying minimal perturbations to the original image to bring a maximum change in the classifier’s prediction and switch the predicted class of the original image [147]. Counterfactual explanation not only helps in identifying the diseased area but also aids in understanding the changes that need to be made to switch the classifier’s prediction. Fig. 3 shows an example of a counterfactual explanation for breast mass prediction in an image patch that is predicted as normal. Counterfactual images are synthesized by using Generative Adversarial Networks (GANs) [46] or by perturbing the latent space of an autoencoder. Generative Adversarial Networks (GANs) [47] are widely used for image synthesis because of their ability to model target data distribution. GANs consist of a generator and a discriminator that work in an opposing manner. GANs can be difficult to train due to loss function instability

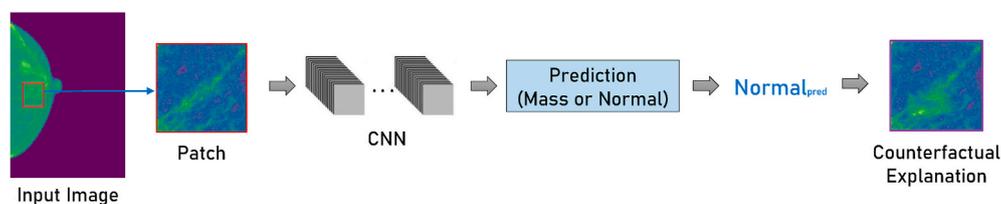
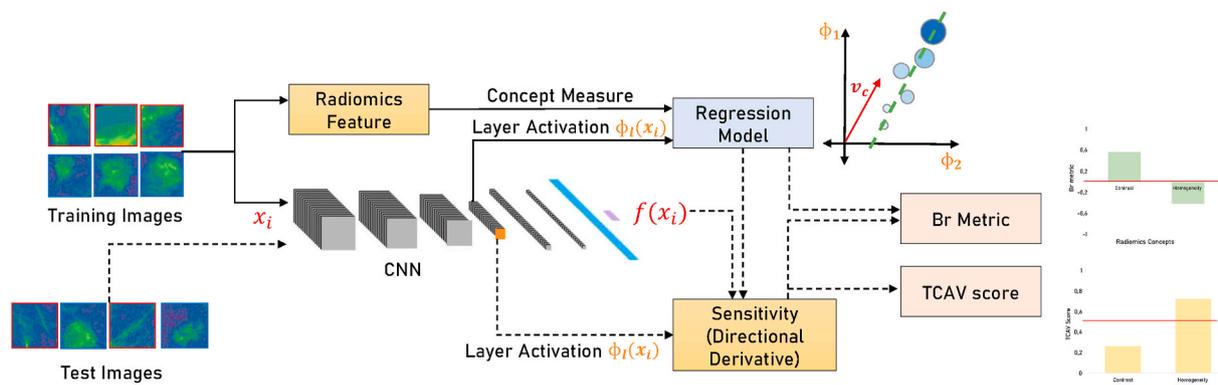


Fig. 3. Counterfactual explanation for breast mass prediction in an image patch from a mammogram. The DL model predicts the label of the input image patch as normal and the counterfactual explanation is obtained by applying minimal perturbation to the input image to change the model’s prediction.

and high sensitivity to hyper-parameters.

Schutte et al. [121] proposed a method using StyleGAN2 [61] to generate synthetic images by a minimal modification that would lead to new predictions. A generator  $G$  has an intermediate latent representation  $z$ . An encoder  $E$  is trained to retrieve the latent representation  $z$  back from the generated image. A classifier  $f$  can be explained by fitting a logistic regression classifier  $f(z_i) = \sigma(\alpha \cdot z_i + \beta)$  in the latent space  $z$  to estimate the labels predicted by  $f$ . The latent representation  $z = E(X)$  of an input image  $X$  is translated in the direction  $\alpha$  to produce a synthetic image ( $X' = G(z + \lambda \cdot \alpha)$ ) that has a lower or higher prediction depending on the value of  $\lambda$ . This method was validated on knee osteoarthritis severity prediction on X-ray images and tumor detection on histology images of metastatic lymph nodes. Narayanaswamy et al. [98] used unsupervised image-to-image translation using CycleGAN [170] to generate counterfactual explanations. Translation between the two classes was applied successively to amplify the differences. A linear SVM trained using features obtained from identified novel image regions achieved performance similar to that of CNNs. This method was demonstrated using retinal images for predicting Diabetic Macular Edema. Ghandeharioun et al. [42] proposed an approach called DISSECT that jointly trains a generator, discriminator, and a concept disentangler to generate a series of examples with increasing degree of concepts that affect the classifier decision. This method was validated on a synthetic dataset that reflected real-world characteristics of melanoma skin lesions. Baumgartner et al. [12] proposed an attribution method based on Wasserstein Generative Adversarial Networks (WGAN) [4] that requires a baseline class  $y$  and a class of interest  $x$ . WGAN estimates a function  $M(x_i)$  such that when it is added to the image from class  $x_i$ , it becomes indistinguishable from class  $y$ . This attribution technique was validated on a synthetic dataset and MRIs from patients with Alzheimer’s disease and mild cognitive impairment. Bass et al. [10] proposed a framework for simultaneous classification and generation of attribution map. This method employed two variational autoencoders (VAE) for generating class relevant and irrelevant latent space for each input image. A discriminator loss on the decoder is used to extract the class irrelevant latent space. The class irrelevant latent spaces of two images from class  $x$  and  $y$  are swapped and fed to a generator to produce an attribution map. The class-relevant latent space is used for the classification and visualization of variations between classes. This method was validated on three brain imaging datasets.

The latent space of an autoencoder can be perturbed to synthesize new images but the resolution of the generated images is limited. Cohen et al. [24] trained an autoencoder with an Encoder  $E(x)$  and a Decoder  $D(z)$  in order to produce a latent representation  $z$  for an input image  $x$ . Perturbations in the latent space for a classifier  $f$  can be produced as:  $z_y = z + \lambda \frac{\partial f(D(z))}{\partial z}$  to produce a new image  $x'_y = D(z_y)$ . Different images  $x'_y$  can be combined for different values of  $\lambda$  to produce a gif. This method was applied to explain chest X-ray classifiers. Seah et al. [122] utilized Generative Visual Rationales (GVRs) for the interpretability of CNNs for predicting congestive heart failure from chest radiographs. A GAN architecture is trained on an unlabeled dataset and it is capable of generating realistic images. An autoencoder is trained to produce latent space representation. The latent space is permuted until the disease is no longer predicted. The generator trained on the unlabeled dataset then



**Fig. 4.** Regression Concept Vectors for global interpretability. A linear regression model is trained to estimate radiomics features using the activations  $\phi_l$  of layer  $l$  from a trained CNN. The influence of radiomics features on the decision of a particular class is obtained by calculating the directional derivation in the direction of increase of radiomics concepts during the testing phase (dashed lines). The effect of radiomics features is quantified in terms of Bidirectional relevance (Br) metric and Testing with Concept Activation Vectors (TCAV) score.

synthesizes a healthy image from the perturbed latent space. The difference of the diseased and reconstructed healthy radiograph produces the GVR.

#### 2.4. Concept attribution

Concept attribution provides global explanations for the deep neural network in terms of high-level image concepts [91]. Testing with Concept Activation Vectors (TCAVs) method quantifies the influence of a high-level image feature on the decision of the model [65]. A linear classifier is trained to differentiate between examples containing the concept of interest and random examples. Concept Activation Vector (CAV) is orthogonal to the classification boundary of this linear classifier. TCAV method utilizes concept activation vector and directional derivative to determine the importance of a particular concept for classification in terms of TCAV score. Kim et al. [65] used TCAV to provide explanations for a model that predicted Diabetic Retinopathy (DR) level using a five-point score. The TCAV score was high for concepts that are relevant for the diagnosis of a particular DR Level but low for concepts that are not associated with a diagnosis of a particular DR level. Gamble et al. [40] used TCAV for interpretability of DL models for breast cancer biomarker status prediction. Clough et al. [23] used a series of fully connected layers that process the latent space of a variational autoencoder (VAE) for cardiac disease classification in CMR images. The sensitivity of the classifier to several clinical biomarkers was evaluated using TCAV in the activation space of a fully connected layer with 64 units. TCAV showed that ventricular ejection and filling rates are the most relevant features for the classifier.

It can be challenging to create a labeled dataset for different concepts to obtain the concept activation vectors especially in the field of medical imaging. Automated Concept-based Explanation (ACE) method mitigates the need for human supervision for labeling the concepts [45]. In ACE, segments are extracted from images of the same class at multiple resolutions. These segments are resized to the original image size. The Euclidean distance in the activation space of the bottleneck layers is used to group similar segments as concepts. Janik et al. [58] used superpixels as segments and consequently employed the ACE method for automated concept discovery. These concepts were then used for TCAV to interpret different cardiac conditions from a DL model for segmentation in Cardiac Magnetic Resonance (CMR).

Radiomics converts medical images into quantitative features that can be utilized for clinical decision support [72,73]. Radiomics features can be used as continuous-valued concepts as they are understandable by humans. Graziani et al. [49] extended TCAV so that continuous variables such as radiomics features could be used as concepts. Regression Concept Vector (RCV) in the direction of greatest increase of

the concept measure is obtained by solving least the square linear regression problem in the activation space of a particular layer. A Bidirectional relevance (Br) metric along with the TCAV score is used to measure the importance of the continuous-valued concept. This approach was validated on breast cancer histopathology classification. Testing with regression concept vectors identified contrast and correlation as the most relevant concepts for the classification task. Fig. 4 presents an overview of how RCV uses radiomics features for global interpretability. Yeche et al. [159] used regression concept vectors to identify important radiomics features for calcification and mass classification from patches extracted from mammograms. A new metric called the Unit Ball Surface Sampling metric (UBS) was introduced to overcome the shortcomings of the Br metric and TCAV score. UBS metric does not depend on the current layer and is indifferent to the concept's representation in feature space.

Couteaux et al. [25] proposed a method to test the robustness and sensitivity of the deep segmentation networks based on radiomics features. The network is robust to a feature if it is not sensitive to its variations. The input image is altered to produce a positive response in the direction of the steepest slope in the input space using network gradients. At each iteration, radiomics features of the segmentation mask are calculated. The network is sensitive to the radiomics features that change during this process of activation maximization. This method was applied to interpret deep neural networks for liver tumor segmentation.

#### 2.5. Language description

Deep neural networks can provide explanations in the form of text along with predictions. The textual descriptions can consist of natural language justifications that are learned in a supervised manner or emergent language justifications that are learned in an unsupervised manner.

##### 2.5.1. Natural language

Deep Neural Networks can be trained to provide textual justifications along with the predictions [149,166]. Structured diagnostic reports along with images can be utilized for the training of such networks. It is challenging to train networks that provide natural language descriptions because of the complexity of the natural language and the limited number of available structured medical reports for training. Lee et al. [76] utilized a justification model on top of the diagnosis network to provide textual and visual reasons for the predictions. To alleviate the problem of the duplication of sentences from the training set, a visual word constraint loss was proposed. This approach was validated for breast mass classification and the textual justifications were evaluated using a sentence dataset that described the shape and margin of the

masses. Zhang et al. [165] proposed an interpretable framework for diagnostic pathology consisting of three cascading networks. The model generating textual justifications is trained using images and associated diagnostic reports. Two experienced pathologists outlined five cellular features in a paragraph. Each paragraph was rephrased by two doctors and two trained students resulting in five paragraphs per image. This method was validated on bladder cancer pathology slides. These methods have added annotated costs because they require the medical reports to be re-formatted in a structured and concise way.

### 2.5.2. Emergent language

Symbolic emergent language is the symbolic representation that is learned spontaneously during communication between two different entities. The generated symbolic expressions can be traced back to the input to provide a layer of interpretability. A symbolic segmentation framework based on emergent language was proposed by Santamaría-Pang et al. [117] for interpretable segmentation. It consists of a U-Net like architecture and two Long Short-Term Memory (LSTM) agents, a Sender and a Receiver. The sender receives a segmentation mask from the U-Net and generates a symbolic sentence. The receiver works in co-operation with the sender and generates the segmentation mask based on the symbolic sentence. Hence, the network learns a language in order to communicate between the two modules. This method was validated on brain tumor segmentation in FLAIR images. The generated symbolic expressions are interpretable in terms of tumor characteristics. Chowdhury et al. [22] used an emergent language-based CNN comprising a sender, symbol generator, and a receiver to perform immune cell marker-based classification and chest X-ray classification. The generated symbols correspond to representative information in the input images.

### 2.6. Latent space interpretation

The latent space of a convolutional neural network refers to the compressed representation of the input image. Disentanglement representation of the latent space aims to model salient factors of variation in the data independently [52]. An autoencoder is used for non-linear dimensionality reduction and it consists of two parts, an encoder, and a decoder. The encoder converts the high-dimensional input image into a compressed representation in the bottleneck layer. The latent space in the bottleneck layer is then used to reconstruct the input image by the decoder. An interpretable latent space aims to capture the salient characteristics of the data that can be understood by humans. Variational Autoencoders (VAE) aim to reconstruct the input by learning a continuous latent representation based on variational Bayesian inference [68]. VAE are generative models that also aim to achieve the disentanglement of latent space.

Puyol-Antón et al. [106] utilized the latent space of VAE for cardiac resynchronization therapy response prediction using fully connected layers. Secondary Classifiers are then trained using a part of latent space to incorporate clinical domain knowledge and disentangle the latent space. It is ensured that secondary classifiers do not use a part of latent space so that it can account for unknown contributing factors. Yang et al. [157] proposed a framework for cross-modality liver segmentation that generated domain agnostic and anatomy preserving disentangled representation. Visualization of the generated domain agnostic and cross-modality images showed that the disentangled representation helps in the interpretation of the model. Chen et al. [21] proposed a method called concept whitening that decorrelates and normalizes the latent space and can be used as an alternative to batch normalization for any layer without loss in performance. Pre-defined concepts are then associated with the axes of the disentangled latent space to provide insight into how these concepts are gradually learned. This method was validated for skin lesion diagnosis. Biffi et al. [13] visualized the anatomical variability encoded in the latent space to make the classification task interpretable. Clough et al. [23] interpolated the latent space

of VAE in the direction of the concept activation vector to visualize the behavior of the classifier. However, the quality of the generated images is not high enough to observe the changes.

The feature space of the CNN consists of hundreds of features even in the bottleneck layers of segmentation networks and fully connected layers of the classification networks. It is not possible to interpret such a high dimensional feature space to visualize intra-class and inter-class similarities, outliers, and other data structure information. Therefore, there is a need to visualize these features in a low-dimensional space such that the significant structure of the data is preserved. Principal Component Analysis (PCA) is a linear technique that utilizes a global covariance matrix to reduce dimensionality and preserve the distances between input samples. Yang et al. [155] used UNet with ResNet blocks for brain ventricle segmentation from multi-modal MRI images. The latent space from the encoder is projected into low dimensional space  $z_l$  using PCA. A multi-layer perceptron model is trained to predict Dice scores using  $z_l$  to estimate the dice scores for the entire input space. 2D projections of  $z_l$  showed that input images with certain characteristic features are clustered together. The relationship among features in a non-linear high dimensional space cannot be expressed in a linear form. t-distributed Stochastic Neighbor Embedding (t-SNE) is a non-linear and iterative dimensionality reduction method [92]. t-SNE first calculates the probability of similarity of the features in the high dimensional space and then in the low dimensional space. Kullback-Leibler divergence is used to minimize the difference in probability distributions using gradient descent. Hence, t-SNE can be used to visualize the patterns in low dimensional space. Dinsdale et al. [30] used t-SNE plots for detecting the bias in the dataset from different scanners by visualizing the features from the fully connected layer on a two-dimensional plot. t-SNE plots were used to validate the proposed iterative approach for unlearning the scanner information to produce scanner invariant features. Fig. 5 shows how t-SNE plots can be used for interpretability to identify dataset bias due to differences in acquisition. Faust et al. [38] utilized t-SNE for generating two-dimensional plots from CNN's final hidden layer in order to visualize the clusters of similar classes. They used t-SNE plots to provide alternate classification output by the k-nearest neighbor approach. Projective Latent Interventions (PLIs) for retraining the classifier by back-propagation based on parametric approximation of t-SNE embedding [53]. This method was validated for standard plane classification for fetal ultrasound imaging.

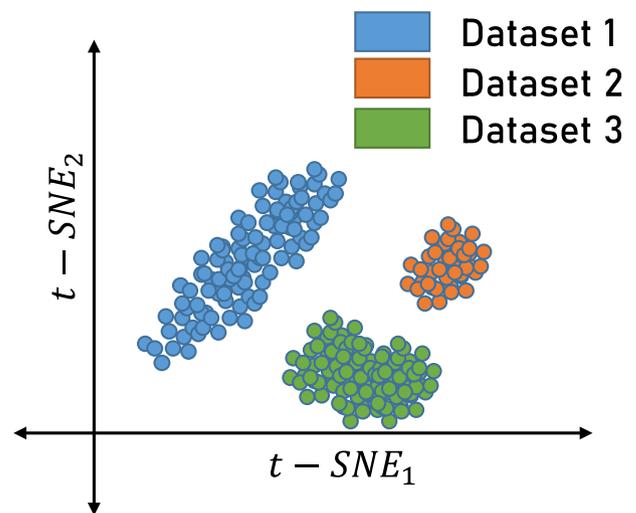


Fig. 5. Low dimensional visualization of a fully connected layer of a CNN using t-distributed stochastic neighbor embedding (t-SNE) for interpretability. This plot shows dataset bias in the DL model due to difference in acquisition.

2.7. Attribution map

The DL models can be explained by highlighting regions of the input image that are relevant for the prediction. These heatmap-based explanations do not offer any information on how these salient regions contribute to the prediction. Fig. 6 shows attribution maps generated by different interpretability methods for explaining a DL model that detects breast cancer in an image patch extracted from a digital mammogram. Table 2 shows examples of different methods related to medical image analysis tasks that use attribution maps for the interpretability of DL models.

2.7.1. Gradient-based methods

Gradient-based methods generate post-hoc attribution maps by utilizing gradients or backpropagation to identify the important parts of the input image for the prediction. Class Activation Maps (CAMs) localize class-specific image regions that the model considers important for classification [168]. To generate CAMs, a global average pooling layer is added after the last convolutional layer. The output of the global average pooling layer is then linearly combined to produce class predictions. CAM for each class is then obtained by taking a weighted sum of the last convolutional layer activations. This interpretability comes at a cost of performance, as the proposed network cannot incorporate fully connected layers. Camalan et al. [16] used CAMs for interpreting the results

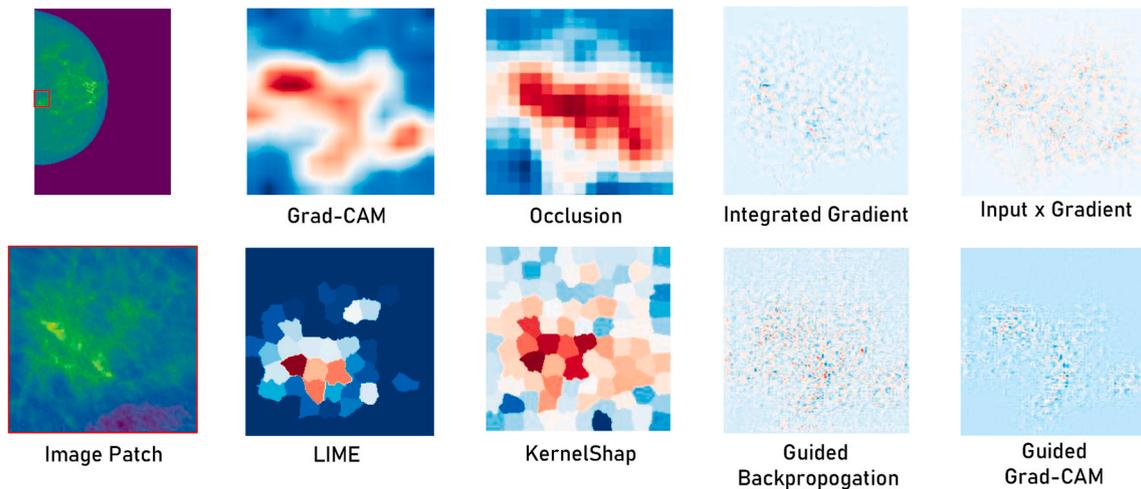


Fig. 6. Attribution maps generated by different interpretability methods for explaining a DL model that detects breast mass in an image patch extracted from a digital mammogram. The first column shows the image patch extracted from the mammogram that is provided to the DL model as input. Red color corresponds to important regions and the blue color corresponds to irrelevant image regions for classification.

Table 2

Examples of different methods related to medical image analysis tasks that use attribution maps for the interpretability of DL models.

Method	Reference	Application
Layerwise Relevance Propagation (LRP)	[15]	Alzheimer’s disease classification
	[35]	Multiple Sclerosis diagnosis
Class Activation Maps (CAM)	[16]	Oral cancer classification
Gradient-Class Activation Maps (Grad-CAM)	[104]	Automatic brain tumor grading
	[101]	Detection of COVID-19 from Chest X-ray and CT scans
Integrated Gradient (IG)	[137]	Diabetic Retinopathy (DR) prediction
	[150]	Multiple Sclerosis classification
Occlusion	[64]	Diagnosis of age-related macular degeneration and diabetic macular edema in OCT images
	[139]	Diagnosis of Alzheimer’s disease
Local Interpretable Model-agnostic Explanations (LIME)	[93]	Parkinson’s disease detection
	[122]	Congestive heart failure prediction
kernel SHAP (Linear LIME + Shapley values)	[161]	Skin cancer detection
	[171]	Lung nodule classification
Trainable Attention	[154]	Melanoma recognition
	[156]	Classification of breast cancer microscopy images
	[59]	Organ segmentation in 3D abdominal CT scans
SmoothGrad	[44]	Identification of cardiac structure, estimation of cardiac function and prediction of systemic phenotypes from Echocardiography
	[102]	Classification of estrogen receptor status from breast MRI
Guided BackPropagation (GBP)	[51]	Lung adenocarcinoma classification
DeepLIFT (Learning Important FeaTures)	[89]	Diagnosis of Multiple Sclerosis
Deep SHAP (DeepLIFT + Shapley values)	[62]	Breast lesion classification, lung lesion classification
	[107]	COVID and Pneumonia classification from chest X-rays
Deep Taylor Decomposition (Deep Taylor)	[127]	Content-based image retrieval for pleural effusion condition in Chest X-Ray images
Multi-Layer Class Activation Maps (MLCAM)	[57]	Feature localization for Confocal Laser Endomicroscopy Glioma images
Expected Gradients	[28]	COVID-19 detection
Contextual Decomposition Explanation Penalization (CDEP)	[109]	Skin Lesion Classification

of the classification model for oral cancer. Multi-Layer Class Activation Maps (MLCAM) is an extension of CAM that can be incorporated at different CNN layers [57]. ML-CAM was used to identify diagnostic features of human gliomas in Confocal Laser Endomicroscopy images.

Gradient-Class Activation Maps (Grad-CAM) produce visual explanations that do not require re-training or changes to the architecture like Class Activation Maps [124]. Grad-CAM allows to explain activations in any layer of the network. The gradient of class  $c$  ( $y^c$  before softmax) is computed with respect to the feature map activations of a particular layer  $A^k$  in order to determine importance  $\alpha_k^c$  of a particular feature map for the target class  $c$  as  $\alpha_k^c = \frac{1}{2} \sum_i \sum_j \frac{\partial y^c}{\partial A^k}$ . The weighted sum of the activations of the layer  $A^k$  with  $\alpha_k^c$  gives the Grad-CAM output  $G^c$  as:  $G^c = \text{ReLU}(\sum_k \alpha_k^c A^k)$ .  $G^c$  is then upsampled to match the input image size. Grad-CAM's output for the last layers in the network lacks fine-grained details. Guided Grad-CAM fuses Guided Backpropagation [134] and Grad-CAM by element-wise multiplication to retrieve the fine details. Grad-CAM++ is an extension of Grad-CAM that provides better localization and is capable of detecting multiple object occurrences [19]. Pereira et al. [104] used Grad-CAM for the interpretability of the DL model for automatic brain tumor grading from MRI data.

The Integrated Gradient (IG) method proposed by Sundararajan et al. [137] mitigates the saturation problem of gradients. In this method, it is critical to select a baseline that corresponds to a near-zero score. A complete black image is a suitable choice for a baseline. The gradients are aggregated for all the points occurring in small steps between the input and baseline. IG method satisfies the desirable attribution axioms of sensitivity and implementation variance. Sundararajan et al. [137] used attribution maps generated by the integrated gradient method to interpret the DL model for Diabetic Retinopathy (DR) prediction. Wargnier-Dauchelle et al. [150] used integrated gradients for the interpretability of a model for Multiple Sclerosis diagnosis. Sayres et al. [118] conducted a collaborative human-AI study to evaluate the effect of using IG attribution maps for DR grading. Ten ophthalmologists performed the task with and without IG explanations. The accuracy of diagnosis increased for cases with DR when the ophthalmologists used attribution maps.

Saliency maps proposed by Simonyan et al. [128] highlight regions of the input image that the deep neural network considers important for prediction by computing the gradient of the output concerning the input pixels using backpropagation. Gradient \* Input has an advantage in that it allows to make use of the input information for better visualization [126]. SmoothGrad is a technique that adds noise to the image to generate new images and produces a final attribution map by taking an average of the attribution map of these similar images [132]. Ghorbani et al. [44] used SmoothGrad with a simple gradient method for interpretability of the DL model applied to echocardiography. Respond-CAM takes a weighted average of the feature maps with their corresponding gradients and generates an attribution map [167]. Respond-CAM demonstrated better performance when compared with Grad-CAM for 3D biomedical images. Saleem et al. [114] proposed an extension of method [135] for generating attribution maps independent of gradients. This method was evaluated qualitatively and quantitatively for the brain tumor segmentation network. DeConvNet consists of a series of deconvolution and unpooling layers and produces an attribution map by setting the negative gradients to zero during the backward pass [162]. Guided BackPropagation (GBP) is a modification of DeConvolution approach [134]. It works by computing the gradient and setting the negative gradients to zero during backpropagation. He et al. [51] used GBP for interpreting decisions made by CNN for lung adenocarcinoma classification. Wickström et al. [152] utilized GBP for interpretability of CNN for colorectal polyps segmentation.

### 2.7.2. Perturbation-based methods

Perturbation-based approaches investigate the effect of altering different parts of the input image on the model's predictions. The

occlusion method produces an attribution map by perturbing the image in a systematic manner to observe the effect on the output. The important parts of the image will have a strong effect on the output when they are altered. Multiple inferences need to be performed for the same input image. It can be computationally expensive to generate occlusion attribution maps if small portions of the image are perturbed. The resolution of the attribution map is constrained by the choice of the patch size that is altered at a time. Occlusion was first utilized by Zeiler and Fergus [162] for the interpretability of deep neural networks. Kermary et al. [64] used occlusion maps for interpretability of a model developed for diagnosing age-related macular degeneration and diabetic macular edema in optical coherence tomography images. Tang et al. [139] utilized occlusion maps for interpretability of a model developed for the diagnosis of Alzheimer's disease.

Local Interpretable Model-agnostic Explanations (LIME) can be used to explain a model by approximating its behavior locally using a surrogate weighted regression model [108]. The training data for LIME is generated by randomly turning a subset of superpixels on or off. Each super pixel corresponds to an input feature and the model's predictions for these perturbed images are the target values. Each perturbed image is weighed by its similarity with the input image using the cosine distance metric. A weighted regression model is trained to estimate the feature importance. Magesh et al. [93] developed a model for the detection of Parkinson's disease in DaTSCANs and utilized LIME for interpretability. Seah et al. [122] performed a comparative analysis of LIME and their proposed attribution method called Generative Visual Rationale for interpretability of a model predicting congestive heart failure in the chest radiograph.

Lundberg and Lee [90] proposed kernel SHAP (Linear LIME + Shapley values) that is similar to LIME except that weights of the regression surrogate model are given by  $\pi_x(z') = \frac{M-1}{\binom{M-1}{|z'|} \cdot (M-|z'|)}$  instead of cosine distance metric. Here,  $M$  represents the number of features and  $|z'|$  represents the number of non-zero features in the input. Young et al. [161] used kernel SHAP and Grad-CAM to show that model with similar performance use completely different rationales for diagnosing skin cancer.

Perturbation-based methods can alter parts of the image that are not understandable in clinical terms. Therefore, there is a need for meaningful perturbations. Zhu and Ogino [171] proposed Guideline-based Additive eXplanation (GAX) that mitigates this problem by first generating features using rule-based segmentation and anatomical irregularities according to the set guidelines. Then, a perturbation-based analysis is performed to obtain understandable explanations in terms of feature importance. GAX method was validated for lung nodule classification. Uzunova et al. [145] generated meaningful perturbations for explaining a classification network for OCT images and brain lesion MRIs by replacing the perturbed area with the healthy equivalent generated using variational autoencoders.

### 2.7.3. Layerwise relevance propagation

Layerwise Relevance Propagation (LRP) is a method based on pixel-wise decomposition of non-linear classifiers that generates a heatmap by evaluating a relevance score [6]. The magnitude of contribution of each neuron in a particular layer is identified in terms of relevance score in a backwards fashion, one layer after the other, starting from the last layer. The relevance score  $R_j^{(l)}$  of the neuron  $j$  in layer  $l$  is calculated from the relevance score  $R_k^{(l+1)}$  of the neuron  $k$  in the layer  $l + 1$  using  $R_j^{(l)} =$

$\sum_k \frac{a_j^{(l)} \times w_{jk}}{a_j^{(l)} \times w_{jk}} \times R_k^{(l+1)}$  where  $a_j^{(l)}$  represents the activation of a neuron  $j$  and  $w_{jk}$  represents the weights between the neurons  $j$  and  $k$ . LRP ensures that the total relevance for all the layers starting from the classification output  $f(x)$  to the input layer is the same. Under these constraints, a neuron is highly relevant if it has a high activation and a high contribution for a neuron of the next layer that have a high

relevance score.

Boehle et al. [15] demonstrated that LRP can be used to explain Alzheimer's disease classification in 3D CNNs with high inter-patient variability and quantitatively showed that LRP relevance maps correlate with clinical knowledge. Eitel et al. [35] evaluated LRP visualizations quantitatively for Multiple Sclerosis (MS) diagnosis and concluded that LRP focuses not only on the individual lesions but also on conventional MRI biomarkers in MS. Thomas et al. [141] utilized LRP to identify physiologically appropriate brain regions corresponding to the cognitive state decoded from fMRI data using recurrent neural networks. Yan et al. [153] used LRP to highlight relevant information in high dimensional fMRI data. Eitel and Ritter [33] carried out a quantitative comparison of four attribution methods for Alzheimer's disease classification and showed that LRP and GBP produce the most coherent explanations.

#### 2.7.4. Trainable attention

Trainable attention modules enable the model to focus on important regions of the image and help to suppress the irrelevant background information during training. Jetley et al. [59] proposed a method for trainable attention that can be incorporated into standard CNN pipelines. This attention mechanism not only improves classification performance but also produces a fine-grained attribution map that focuses on key parts of the image. The attention module calculates a compatibility score  $c_i^s$  between the local features  $f_i^s$  obtained from the convolutional layer  $s$  where  $i$  represents the spatial location within  $s$  and the global features  $g$  obtained from the final convolutional layer. The attention weights  $\alpha_i^s$  can be obtained from the compatibility scores using a softmax function:  $\alpha_i^s = \frac{\exp(c_i^s)}{\sum_j \exp(c_j^s)}$ . The final output of the attention module  $g_a^s$  is calculated by taking a weighted sum of local features using attention weights as  $g_a^s = \sum_{i=1}^n \alpha_i^s \cdot f_i^s$ . The output of the attention modules  $g_a$  from various layers are then concatenated and fed into a fully connected layer to perform classification. The attention module tries to ensure that the compatibility score is high only when the image patch contains the part relevant for classification. The greatest benefit is achieved by using these attention modules relatively late in the CNN pipeline. Yan et al. [154] incorporated attention modules in the network to generate attention maps that highlighted image regions relevant for melanoma recognition. The addition of the attention modules resulted in a performance improvement. Li et al. [79] proposed Guided Attention Inference Network (GAIN) for generating attention maps for weakly supervised tasks by incorporating attention mining loss and a method for improving these attention maps by utilizing extra supervision when available. Yang et al. [156] utilized this guided soft-attention mechanism for the classification of breast cancer microscopy images.

Schlemper et al. [119] proposed attention gates that can be incorporated into standard CNN pipelines. These attention gates are more specific to local regions and achieve superior performance compared to global feature vectors utilized in Ref. [59]. The proposed attention gates can also be utilized in segmentation networks unlike in Ref. [59]. The use of attention gates demonstrates superior classification performance in scan plane detection for the screening of fetal ultrasound. A gain in performance was achieved by incorporating attention gates in a

standard U-Net for the segmentation of multiple organs in the 3D CT abdominal dataset. The fine-grained attention maps can be visualized to interpret the predictions. Li et al. [81] used attention gates in Dense-U-Net for breast mass segmentation in mammograms. Sun et al. [136] used spatial and channel-wise attention in the decoder of the proposed U-Net architecture to visualize features learned at every resolution.

#### 2.7.5. Other methods

Shrikumar et al. [126] proposed DeepLIFT (Learning Important Features) explanation method that assigns importance to the neurons based on the activation of each neuron compared to its reference activation [128,134,162]. can produce misleading results when the local gradient is zero. DeepLIFT does not face this problem because the difference from the reference value may be non-zero even when the local gradient is zero. Lundberg and Lee [90] proposed Deep SHAP that approximates SHAP values by adapting DeepLIFT. Montavon et al. [97] proposed an attribution method called Deep Taylor Decomposition (Deep Taylor) that applies Taylor decomposition in a layer-wise manner. Silva et al. [127] utilized the attribution maps generated by the Deep Taylor method for an automatic content-based medical image retrieval system. Prior knowledge can be utilized to constrain explanations during training to improve model performance and produce high-quality attribution maps [37].

Rieger et al. [109] introduced Contextual Decomposition Explanation Penalization (CDEP) to incorporate explanation error in the loss function and demonstrated that the Grad-CAM heatmaps with CDEP regularization ignore confounding variables during skin lesion classification. DeGrave et al. [28] used Expected Gradients method [37] that utilizes attribution prior for generating attribution maps with desirable properties like smoothness and sparsity for the detection of COVID-19 from chest radiographs. Kindermans et al. [67] showed that LRP, Guided-BackProp, and DeConvNet do not produce theoretically correct explanations for single-layer architecture i.e. linear models. They proposed PatternNet and PatternAttribution explanation methods that provide correct signal visualization for linear models.

#### 2.8. Anatomical prior

Deep Neural Networks for medical image analysis problems can be made interpretable by incorporating task-specific structural information/anatomical prior in the workflow. Midline shift (MLS) in the brain is an important characteristic feature that can be used for the diagnosis of traumatic brain injury, brain tumors and some other brain abnormalities [85]. Fig. 7 shows how anatomical prior is incorporated in the DL model for midline shift detection. A simple approach for predicting MLS could employ a CNN for direct diagnosis using MRI slices. It would be difficult to interpret the predictions of such a model as several key points need to be identified for MLS prediction and their relative information is important. Pisov et al. [105] proposed an interpretable approach by first estimating the mid-line exploiting the structural knowledge and consequently predicting Mid-Line Shift (MLS) in a brain MRI from the estimated curve. Segmentation models can be trained to localize the region of interest so that the network learns only from the

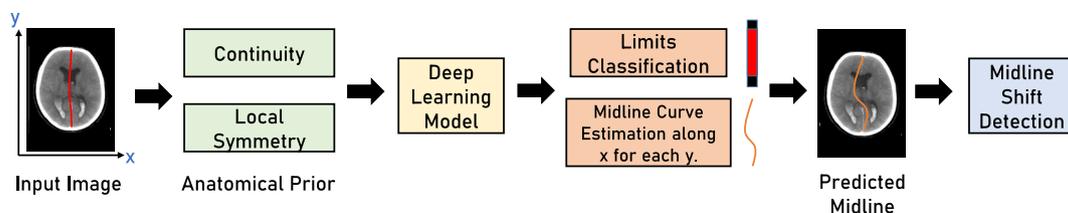


Fig. 7. Continuity and local symmetry of the midline is incorporated in the DL pipeline as an anatomical prior for midline shift detection [105]. This DL solution is more interpretable than a DL model that predicts midline shift directly from the image.

**Table 3**

The table summarizes different methods for the interpretability of DL models for medical image segmentation.

Interpretability Method Category	Reference	Details	Application
Attribution Map	[114]	Visual explanations that are independent of gradients for 3D segmentation networks.	Multi-Modal MRI Brain Tumor Segmentation
	[152]	Guided Backpropagation [134] method for identifying important pixels.	Colorectal Polyps Segmentation
	[59]	Trainable soft attention mechanism that highlights salient parts of the image.	Organ Segmentation in 3D abdominal CT scans
Concept Attribution	[25]	Radiomics features are used to evaluate the robustness and sensitivity of deep segmentation networks.	Liver Tumor Segmentation from CT scans
	[58]	Automated concept discovery [45] and superpixels are used to employ the TCAV method [65] for providing global explanations to interpret different cardiac conditions.	Cardiac MRI Segmentation
Language Description	[117]	Unsupervised symbolic emergent language explanations that are generated using two LSTM agents.	Brain Tumor Segmentation using FLAIR images
Internal network Representation	[99]	Network Dissection [11] and Activation Maximization [36] methods for understanding the concepts learned by different filters in the CNN.	Multi-modal MRI Brain Tumor Segmentation
Latent Space Interpretation	[155]	Low dimensional representation of the latent space for identifying clusters of cases with similar performance.	Multi-modal MRI Brain Ventricle Segmentation
	[157]	Disentangled latent space representation for the interpretability of the segmentation network.	Cross Modality Liver Segmentation

specified region to enhance interpretability. The interpretability of the trained segmentation network then also comes into question. CNNs are more inclined to learn texture features and introducing a shape bias can improve robustness [41]. Anatomical shape prior can also enhance the interpretability of segmentation networks. Sun et al. [136] introduced Shape Attentive U-Net that includes a secondary shape stream parallel to the original image stream in the proposed architecture and generates a shape attention map that can be used for interpretability.

### 2.9. Internal network representation

It can be beneficial to visualize the different features being learned by the filters in a CNN to understand the features the network is using to arrive at a particular decision. These visualizations are meaningful when dealing with natural images, as different filters can represent peculiar structures and natural appearances. However, the structures and patterns in medical images are not very obvious and hard to interpret. Activation Maximization detects the patterns that trigger a particular neuron in a deep network by transforming the pixels of a randomly initialized image to maximize the activation of the neuron keeping the weights of the network and desired output constant [36]. Network Dissection is an interpretability method that quantifies the activation and concept alignment by measuring the intersection over union score between the pixel-wise concept annotations and the corresponding activation of the neuron [11]. Natekar et al. [99] explored the Network Dissection interpretability method to understand information organization in the DL model for brain tumor segmentation. They also used Activation Maximization with regularization to identify different concepts learned by the model.

## 3. Evaluation of interpretability methods

As evident in the previous section, several efforts are being made to explore interpretability methods for understanding the inherent black-box nature of the DL algorithms. It is difficult to determine which method works better for understanding the deep neural networks for a particular medical imaging application, as explanations can be subjective. Explanations also depend on a particular application and the context of their utility. The uncertainty of DL models is also translated to their interpretability, and there is no ground truth available for explanations. All these characteristics make the evaluation and quantification of the quality of interpretability methods a great challenge. Doshi-Velez and Kim [31] proposed three different categories for the evaluation of interpretability methods. Application-grounded evaluations involve experts for a specified application e.g. doctors for diagnosis.

Human-grounded evaluations involve lay humans to test the general quality of explanations. Finally, functionality-grounded evaluations involve proxy tasks instead of humans for evaluating the quality of explanations. Functionality-grounded evaluation metrics that do not involve human interaction are desirable for interpretability due to time and cost constraints. Table 4 shows examples of different evaluation settings used to ensure that the explanations produced by various interpretability methods for medical image analysis tasks are trustworthy.

### 3.1. Evaluation in a clinical setting

Application-grounded evaluation in the presence of a medical doctor is necessary to determine if the end-users are satisfied with the provided explanations. System Causability Scale (SCS) assesses the quality of the interpretability methods which can be utilized in the medical domain [55]. It is also important to determine if the provided explanations introduce a bias towards a specific class that can result in a performance drop. Sayres et al. [118] studied the impact of the DL algorithm and integrated gradient explanations on physicians for diagnosing Diabetic Retinopathy (DR) in a computer-assisted setting. When integrated gradient explanations were used along with model predictions by ten ophthalmologists, it resulted in an over-diagnosis of the disease in normal cases without DR. This may be because the attribution maps generated by the integrated gradient method provide evidence for the occurrence of the disease and they may not be useful when there is no disease present. The human-AI collaboration that utilized explanations achieved overall better accuracy. Therefore, there is a need to evaluate the attribution methods in a clinical setting to determine if the explanations are useful for the clinicians and they are not introducing bias towards one class.

### 3.2. Evaluation using imaging biomarkers

An imaging biomarker can be described as any quantifiable feature or structure in an image that can be used for the diagnosis and prognosis of a disease [151]. DL models automatically detect patterns from the data and learn complex functions that are not understandable to humans. These functions may learn to utilize established biomarkers from the data along with unknown imaging biomarkers for diagnosis. The explanations of DL models can be validated by using known imaging biomarkers to show that the DL model is using clinically relevant features for diagnosis. This validation can foster trust in the utility of DL models in clinical practice. Eitel et al. [35] utilized Layer-wise Relevance Propagation to develop an association between the relevant features that

**Table 4**

Examples of different evaluation strategies used to ensure that the explanations produced by various interpretability methods are trustworthy.

Evaluation	Reference	Details
Evaluation in a Clinical Setting	[55] [118]	System Causability Scale (SCS) can be used by medical doctors to assess the quality of the interpretability methods. Ten ophthalmologists performed Diabetic Retinopathy (DR) grading in three scenarios: using AI model with integrated gradient explanations, using only AI model and in an unassisted setting.
Evaluation using Imaging Biomarkers	[35] [15]	LRP attribution maps for Multiple Sclerosis prediction were associated with established biomarkers. LRP attribution maps for Alzheimer's disease prediction correlate with hippocampal volume which is a known biomarker of the disease.
Evaluation against Adversarial Perturbations	[43]	Input images with small perturbations that are visually indistinguishable produce strikingly different attribution maps using DeepLIFT and Integrated Gradient method.
Quantitative Evaluation of Attribution Maps	[1] [161] [160] [60] [163] [142]	Grad-CAM and Gradient method pass the model parameter randomization test and data randomization test. Reproducibility, model dependence, and sensitivity tests for kernel SHAP and Grad-CAM attribution maps for skin cancer classification. Retraining the model and applying slight translations to assess attribution maps for explaining DL model for osteoporotic fracture discrimination in CT scans, Modality Importance and Modality Specific Feature Importance metrics were used for the evaluation of 16 methods that generate attribution maps for a brain tumor classification task. Area Over Perturbation Curve (AOPC) metric [115] was used for the quantitative evaluation of attributions maps for interpretability of a CNN for fetal head circumference estimation from ultrasound images. A synthetic dataset comprising of cell images with easily distinguishable features was used to quantify and compare attribution maps.
Quantitative Evaluation of Counterfactual Explanation	[131]	Frechet inception distance, counterfactual validity, and foreign object preservation metrics along with two other clinical metrics were used for quantitative evaluation of explanations for chest X-ray classification task.
Quantitative Evaluation of Language Description	[77]	BLEU [103], ROUGE-L [86] and CIDEr [146] metrics were used for the evaluation of generated natural language descriptions against ground truth reference sentences provided by experts.

the model uses and the established biomarkers for Multiple Sclerosis. Imaging biomarkers can be used as concepts to quantify their influence on the model's prediction using concept activation vector analysis [23]. Boehle et al. [15] used LRP to explain deep neural networks for the diagnosis of Alzheimer's disease. It was quantitatively shown that areas that have high relevance correlate well with hippocampal volume, which has been identified as a key biomarker of Alzheimer's disease. High variability in heatmaps between different disease cases showed that LRP might be used to identify biomarkers for different stages of Alzheimer's disease. A cascade of two autoencoders was employed by Waldstein et al. [148] to obtain a global representation of 3D Optical Coherence Tomography (OCT) in just 20 numbers which represented global and local intrinsic features. This low dimensional feature representation was then validated by comparison with the heatmaps of conventional biomarkers obtained by automated image segmentation algorithms as well as using correlation and machine learning regression. Imaging biomarkers can also be utilized for both qualitative and quantitative evaluation of interpretability methods.

### 3.3. Evaluation against adversarial perturbations

Adversarial examples have demonstrated the vulnerabilities present in the state-of-the-art DL solutions [48]. The security of an automated health care solution for diagnosis, treatment planning, and treatment follow-up is critical and should not be compromised by input perturbations [39]. Interpretability methods are also susceptible to similar adversarial attacks. Ghorbani et al. [43] tested the robustness of DeepLIFT and Integrated Gradient attribution maps to small perturbations of input data. The attribution maps of visually indistinguishable input images with similar labels are radically different. Therefore, there is a need to test the robustness of interpretability methods to input perturbations and noise.

### 3.4. Quantitative evaluation for trustworthy explanations

Explanations need to be robust, sensitive to model and data, and consistent. The qualitative evaluation alone cannot ensure these properties for some interpretability methods. There is a need to evaluate explanations using quantitative metrics to ensure that the interpretability methods hold these desirable properties. Different metrics and evaluation strategies have been proposed to overcome the lack of explanation ground truth. It is not possible to come up with quantitative

metrics that are applicable for all interpretability methods [169].

#### 3.4.1. Attribution maps

Attribution maps offer explanation by showing the part of the input image that the model considers important for making the decision. It does not explain how the model is utilizing the relevant information. The explanation heat map for multiple classes may be the same. The correct image area may be highlighted in the heat map even if the prediction is wrong. This can instill a false sense of confidence in the model which can be detrimental for troubleshooting and model deployment in clinical practice [110]. The robustness of the attribution maps produced by different methods cannot be solely evaluated qualitatively. Adebayo et al. [1] utilized model parameter randomization test and data randomization test to quantitatively evaluate if the explanations generated by attribution methods are faithful to the model and the data. Guided BackProp and Guided Grad-CAM fail both the tests while Integrated Gradient and Input \* Gradient fail the data randomization tests. Only Gradient and Grad-CAM pass both the tests. Young et al. [161] performed reproducibility, model dependence, and sensitivity tests for kernel SHAP and Grad CAM attribution maps for skin cancer classification. DL models with similar accuracy can generate strikingly different attribution maps [160,161]. In addition to the tests proposed in Refs. [1,161], Yilmaz et al. [160] assessed the attribution maps for explaining the DL model for osteoporotic fracture discrimination in CT scans by retraining a given model configuration and applying slight translations. It showed that some of the attribution methods behave as edge detectors. Eitel and Ritter [34] trained DL models for Alzheimer's disease classification in identical settings to test the robustness of the attribution maps produced by LRP, Gradient \* Input, GBP, and occlusion methods. SmoothGrad is sensitive to sample size used for taking the average and LIME is sensitive to the number of perturbed instances as well as the random seeds for superpixels generation [7]. Zhang et al. [163] adapted area over perturbation curve (AOPC) metric [115] for the quantitative evaluation of attributions maps for interpretability of a CNN for fetal head circumference estimation from ultrasound images. Attribution maps produced by Gradient and DeConvNet methods were found to be insensitive to the model. Tjoa and Guan [142] utilized a synthetic dataset comprising cell images with easily distinguishable features to quantify and compare attribution maps for the explanation of deep neural networks. A quantitative and qualitative evaluation showed that some methods produce attribution maps that are better at localization while some others at edge detection. It is also important to

evaluate the performance of attribution maps on multi-modal data. Jin et al. [60] utilized Modality Importance and Modality Specific Feature Importance metrics for the evaluation of 16 saliency maps along with a doctor user study on a brain tumor classification task. Most of the attribution maps were not able to localize modality-specific key features.

### 3.4.2. Other methods

When using the TCAV method, a statistical test for a particular concept by training multiple times ensures the stability of the concept. A two-sided  $t$ -test of the TCAV scores should be used to evaluate if the null hypothesis is rejected [65]. Trustworthy Explainability Acceptance metric quantitatively measures the distance between explanations produced by the AI system and the explanations provided by medical experts [63]. Singla et al. [131] utilized three metrics for the evaluation of the counterfactual explanations for chest X-ray classification: Fréchet Inception Distance (FID), Counterfactual Validity (CV), and Foreign Object Preservation (FOP). FID metric assesses visual quality, CV metric quantifies if the counterfactual is compatible with classifier's prediction and FOP metric checks if the patient-specific information is retained. They also used clinical metrics namely cardiothoracic ratio and the score for the detection of normal costophrenic recess to demonstrate the clinical utility of the explanations. Lee et al. [77] used BLEU [103], ROUGE-L [86] and CIDEr [146] metrics for the evaluation of generated natural language descriptions against ground truth reference sentences provided by experts.

## 4. Discussion

Transparency of deep neural networks is an essential clinical, legal, and ethical requirement. We have identified nine different categories of interpretability methods for DL methods. These interpretability methods and their application to medical image analysis problems have been discussed in detail. Evaluation of the interpretability methods is also a critical step to validate the utility of the generated explanations and develop insights. Multiple interpretability methods can be combined to understand the predictions of DL models. Fig. 8 shows that attribution maps and concept attribution post-hoc interpretability methods are used together to provide local and global explanations respectively for mass detection in mammograms. Interpretability of DL networks for segmentation is a difficult problem because attribution of a single-pixel holds little significance. Hence, there are comparatively few interpretability methods available for segmentation networks. Table 3 shows the various interpretability methods for understanding deep neural networks for medical image segmentation.

Post-hoc interpretability methods are only approximations and do not accurately depict the correct model behavior and therefore compromise trust in the explanations [5,110]. An additional model for interpretability can lead to an overly complicated decision system because two models may require troubleshooting instead of one. Since

post-hoc explanation methods are only approximations, they cannot ensure complete accountability [5]. There is a need to evaluate the post-hoc explanations carefully before they can be used in the clinical workflow. Fig. 9 shows different post-hoc interpretability methods. Interpretability can be introduced during the design process of the deep learning models and can help overcome the pitfalls of post-hoc interpretability. Fig. 10 shows different methods that incorporate interpretability during the design process. Domain expertise and additional data may be required to design inherently interpretable models.

### 4.1. Performance and interpretability trade-off

It is incorrectly assumed that the gain in performance of DL models comes at a cost of interpretability [155]. Concept learning models and case based models, which are interpretable by design, have achieved performance at par with black-box models in medical imaging applications. This has debunked the myth of compromise between performance and interpretability. It depends on the researcher's ability to discover the patterns in an interpretable way while at the same time allowing for flexibility to fit the data accurately [110].

### 4.2. Discovery of imaging biomarkers

Explainability methods can hypothesize new imaging biomarkers associated with various diseases. Seegerer et al. [123] utilized Layer-wise Relevance Propagation for understanding the intrinsic details of the model. This revealed several possible imaging biomarkers for the prediction of Estrogen Receptor status from Haematoxylin-Eosin Images i.e. nuclear and stromal morphology and lymphocyte infiltration. Zhang et al. [164] introduced a novel architecture for detecting potential imaging biomarkers when only image labels are available. An encoder-decoder architecture was utilized for generating fake normal images to train the discriminator. The discriminator differentiated between normal and abnormal images. At the same time, the difference of encoder-decoder input and output was fed to a CNN classifier. This difference image identified the potential biomarkers. The approach was validated on diabetic retinopathy and skin images with real and stimulated biomarkers respectively. Several efforts have been made to detect imaging biomarkers for Autism spectrum disorder using interpretability methods for DL models [82,83,172]. To find an imaging biomarker for Alzheimer's disease from molecular imaging data, Ding et al. [29] utilized saliency maps for the interpretation of the DL model. The analysis did not reveal any human interpretable imaging biomarker as the model considered the entire brain volume for the diagnosis. This affirms the need for further improving interpretability methods so that we can better understand the DL models to detect new imaging biomarkers. These imaging biomarkers after validation can then be used for risk assessment and diagnosis in clinical practice. The skin surrounding a lesion has also been identified as a potential biomarker for skin lesion

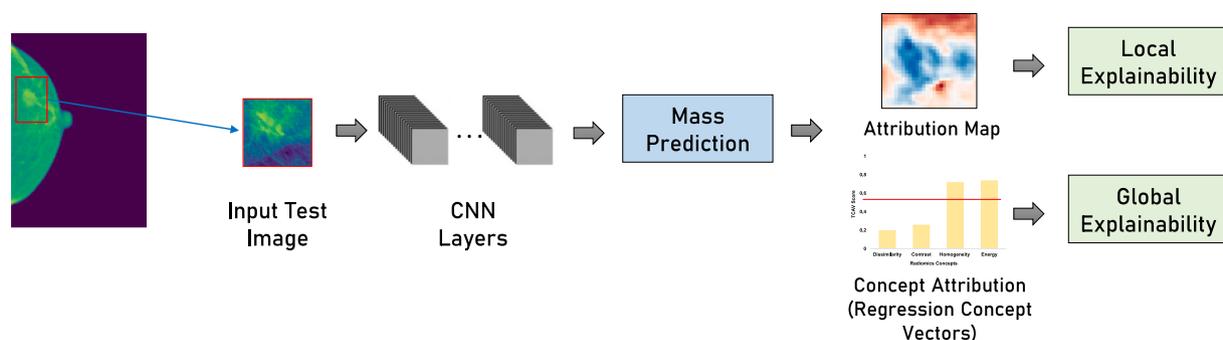


Fig. 8. Post-hoc interpretability for mass detection in mammograms. Attribution map explains each test image by highlighting the regions of the image that the model considers important and testing with regression concept vector method provides global interpretability by quantifying the influence of radiomics features on the prediction for each class.

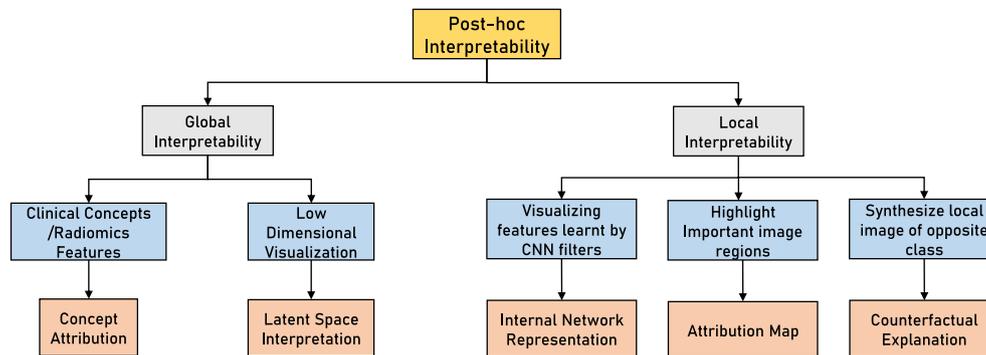


Fig. 9. The flowchart shows different methods for post-hoc interpretability for understanding deep learning models for medical image analysis tasks.

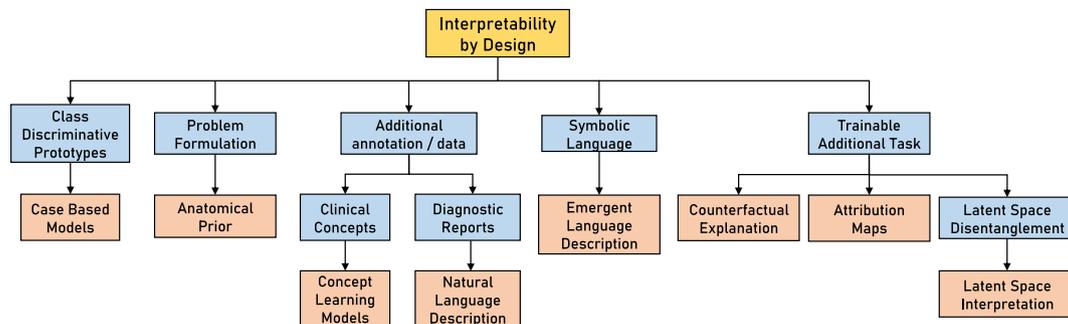


Fig. 10. The flowchart shows different methods that incorporate interpretability during the design process of the deep neural network.

diagnosis using interpretability methods for DL models [84,144]. A study on the comparison of diagnostic performance for skin cancer recognition revealed that AI algorithms achieve better accuracy than human experts [143]. This improvement was particularly prominent for the diagnosis of pigmented actinic keratoses. Tschandl et al. [144] demonstrated that Grad-CAM shows that the CNN pays higher attention to the outside region of the lesion as compared to other categories. This observation is validated by the fact that chronic sun damage is responsible for actinic keratoses and is visible in the surrounding skin region. Medical students were instructed to pay attention to chronic sun damage for diagnosis and this resulted in an overall increase in performance for all categories.

#### 4.3. Validation of explanations using multi-modal data

Validation of explanations can be done using genomics and pathology. Panth et al. [100] investigated the causal relationship between genetic features and radiomics features by designing an experiment to observe the effect of genetic induction and irradiation. Image features can be correlated with tumor biology using ground truth pathology substrates [116]. [50] studied the association between tumor biology and radiomics features using two independent cohorts with lung cancer. Similar experiments need to be performed to validate explanations for DL models.

#### 4.4. Guidelines for using interpretability methods

Interpretability of the deep neural networks should be considered as important as the performance. To foster trust in the DL solutions for medical image analysis tasks, it is important to involve clinicians in the model design process so that the algorithm becomes interpretable inherently e.g. using multi-modal data (text, images), annotating clinical concepts for diagnosis. Both quantitative and qualitative evaluation of post-hoc interpretability methods should ensure the robustness and faithfulness of the explanations. The insights and explanations provided

by the interpretability methods can help discover new imaging biomarkers. Applications-grounded evaluations should be carried out in presence of a clinician to ensure the utility of the explanations and eliminate concerns related to bias.

#### 4.5. Future directions

Case-based models and concept learning models are promising interpretability methods that have the potential to be incorporated into clinical workflow in the future, as they are interpretable by design and they demonstrate performance that is comparable to black-box CNNs. Sanity checks for attribution maps are important to ensure their robustness. Multiple post-hoc interpretability methods can be utilized to understand how the regions highlighted by attribution maps contribute to the final prediction. Multi-modal data consisting of images, texts, and genomics data can complement information to increase performance and enhance interpretability. Holzinger et al. [56] proposed that Graph Neural Networks might be utilized to make use of multi-modal data and knowledge bases to build a human-in-the-loop interactive explainability solution. These counterfactual graphs can establish causal links between multi-modal feature representations that may comprise images, text, genomics data, etc. using graph structures. There is a need to combine human-centered evaluations and quantitative functionality-based evaluations to suggest the most suitable interpretability methods [169].

## 5. Conclusion

The incorporation of deep neural networks in the clinical workflow for medical image analysis tasks is impeded by the vague understanding of their decision-making process. This review paper summarizes the technical details, limitations, and applications of interpretability methods that are available for achieving some degree of transparency for deep neural networks for medical image analysis tasks. Quantitative and qualitative evaluations are necessary to ensure robust and trustworthy explanations. Application-grounded evaluations in the presence

of a medical doctor determine the utility of the interpretability methods and help in detecting if the explanations introduce bias that can result in over-diagnosis or under-diagnosis. Interpretability methods can be validated using imaging biomarkers and they can aid in discovering new imaging biomarkers. Case-based models and concept learning models that are inherently interpretable DL models have achieved performance at par with black-box networks. Post-hoc interpretability methods should be utilized with care as they approximate the model behavior and can instill a false sense of confidence in the AI solution.

### CRedit authorship contribution statement

**Zohaib Salahuddin:** Conceptualization, Methodology, Investigation, Writing – original draft, preparation, Writing – review & editing. **Henry C. Woodruff:** Conceptualization, Methodology, Investigation, Writing – original draft, preparation, Writing – review & editing, Supervision. **Avishek Chatterjee:** Writing – original draft, preparation, Writing – review & editing. **Philippe Lambin:** Writing – review & editing, Supervision, Funding acquisition.

### Declaration of competing interest

We declare that this manuscript is original, has not been published before, and is not currently being considered for publication elsewhere. Dr. Lambin is co-inventor of two non-issued, non-licensed patents on Deep Learning-Radiomics-Histomics (N2024482, N2024889). Dr. Woodruff and Dr. Lambin have (minority) shares in the company Radiomics. The rest of the authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper. All authors have approved the manuscript and agree with its submission to the Journal of Computers in Biology and Medicine. Thank you very much for your kind consideration.

### Acknowledgements

Authors acknowledge financial support from ERC advanced grant (ERC-ADG-2015 n° 694 812 - Hypoximmuno), the European Union's Horizon 2020 research and innovation programme under grant agreement: MSCA-ITN-PREDICT n° 766 276, CHAIMELEON n° 952 172, EuCanImage n° 952 103, TRANSCAN Joint Transnational Call 2016 (JTC2016 CLEARLY n° UM 1757-8295), H2020-JTI-IMI2-2020-23-two-stage and IMI-OPTIMA n° 10 103 434.

### References

- [1] J. Adebayo, J. Gilmer, M. Muelly, I. Goodfellow, M. Hardt, B. Kim, Sanity checks for saliency maps, in: S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, R. Garnett (Eds.), *Advances in Neural Information Processing Systems*, Curran Associates, Inc, 2018.
- [2] H.J.W.L. Aerts, E.R. Velazquez, R.T.H. Leijenaar, C. Parmar, P. Grossmann, S. Carvalho, J. Bussink, R. Monshouwer, B. Haibe-Kains, D. Rietveld, F. Hoebbers, M.M. Rietbergen, C.R. Leemans, A. Dekker, J. Quackenbush, R.J. Gillies, P. Lambin, Decoding tumour phenotype by noninvasive imaging using a quantitative radiomics approach, *Nat. Commun.* 5 (2014) 4006, <https://doi.org/10.1038/ncomms5006>.
- [3] G. Aresta, T. Araújo, S. Kwok, S.S. Chennamsetty, M. Safwan, V. Alex, B. Marami, M. Prastawa, M. Chan, M. Donovan, G. Fernandez, J. Zeineh, M. Kohl, C. Walz, F. Ludwig, S. Braunewell, M. Baust, Q.D. Vu, M.N.N. To, E. Kim, J.T. Kwak, S. Galal, V. Sanchez-Freire, N. Brancati, M. Frucci, D. Riccio, Y. Wang, L. Sun, K. Ma, J. Fang, I. Kone, L. Boulmane, A. Campilho, C. Eloy, A. Polónia, P. Aguiar, Bach: grand challenge on breast cancer histology images, *Med. Image Anal.* 56 (2019) 122–139, <https://doi.org/10.1016/j.media.2019.05.010>. <https://www.sciencedirect.com/science/article/pii/S1361841518307941>.
- [4] M. Arjovsky, S. Chintala, L. Bottou, Wasserstein generative adversarial networks, in: *Proceedings of the 34th International Conference on Machine Learning*, PMLR, 2017, pp. 214–223.
- [5] B. Babic, S. Gerke, T. Evgeniou, I.G. Cohen, Beware explanations from ai in health care, *Science* 373 (2021) 284–286. <https://science.sciencemag.org/content/373/6552/284>.
- [6] S. Bach, A. Binder, G. Montavon, F. Klauschen, K. Müller, W. Samek, On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation, *PLoS One* 10 (2015).
- [7] N. Bansal, C. Agarwal, A.M. Nguyen, Sam: the sensitivity of attribution methods to hyperparameters, in: *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2020, pp. 11–21.
- [8] A.J. Barnett, F.R. Schwartz, C. Tao, C. Chen, Y. Ren, J.Y. Lo, C. Rudin, Interpretable Mammographic Image Classification Using Cased-Based Reasoning and Deep Learning, 2021.
- [9] A. Barredo Arrieta, N. Díaz-Rodríguez, J. Del Ser, A. Bennetot, S. Tabik, A. Barbado, S. Garcia, S. Gil-Lopez, D. Molina, R. Benjamins, R. Chatila, F. Herrera, Explainable artificial intelligence (xai): concepts, taxonomies, opportunities and challenges toward responsible ai, *Inf. Fusion* 58 (2020) 82–115, <https://doi.org/10.1016/j.inffus.2019.12.012>. <https://www.sciencedirect.com/science/article/pii/S1566253519308103>.
- [10] C. Bass, M. da Silva, C. Sudre, P.D. Tudosiu, S. Smith, E. Robinson, Icam: interpretable classification via disentangled representations and feature attribution mapping, in: H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, H. Lin (Eds.), *Advances in Neural Information Processing Systems*, Curran Associates, Inc., 2020, pp. 7697–7709.
- [11] D. Bau, B. Zhou, A. Khosla, A. Oliva, A. Torralba, Network dissection: quantifying interpretability of deep visual representations, in: *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 3319–3327.
- [12] C.F. Baumgartner, L.M. Koch, K.C. Tezcan, J.X. Ang, E. Konukoglu, Visual feature attribution using wasserstein gans, in: *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018, pp. 8309–8319.
- [13] C. Biffi, J. Cerrolaza, G. Tarroni, W. Bai, A. de Marvao, O. Oktay, C. Ledig, L. L. Folgoc, K. Kamnitsas, G. Doumou, J. Duan, S. Prasad, S. Cook, D. O'Regan, D. Rueckert, Explainable anatomical shape analysis through deep hierarchical generative models, *IEEE Trans. Med. Imag.* 39 (2020) 2088–2099.
- [14] P. Bándi, O. Geessink, Q. Manson, M. Van Dijk, M. Balkenhol, M. Hermsen, B. Ehteshami Bejnordi, B. Lee, K. Paeng, A. Zhong, Q. Li, F.G. Zanjani, S. Zinger, K. Fukuta, D. Komura, V. Ovtcharov, S. Cheng, S. Zeng, J. Thagaard, A.B. Dahl, H. Lin, H. Chen, L. Jacobsson, M. Hedlund, M. Çetin, E. Halici, H. Jackson, R. Chen, F. Both, J. Franke, H. Küsters-Vandeveld, W. Vreuls, P. Bult, B. van Ginneken, J. van der Laak, G. Litjens, From detection of individual metastases to classification of lymph node status at the patient level: the camelyon17 challenge, *IEEE Trans. Med. Imag.* 38 (2019) 550–560, <https://doi.org/10.1109/TMI.2018.2867350>.
- [15] M.D. Boehle, F. Eitel, M. Weygandt, K. Ritter, Layer-wise relevance propagation for explaining deep neural network decisions in mri-based alzheimer's disease classification, *Front. Aging Neurosci.* 11 (2019).
- [16] S. Camalan, H. Mahmood, H. Binol, A.L.D. Araújo, A.R. Santos-Silva, P.A. Vargas, M.A. Lopes, S.A. Khurram, M.N. Gurcan, Convolutional neural network-based clinical predictors of oral dysplasia: class activation map analysis of deep learning results, *Cancers* 13 (2021), <https://doi.org/10.3390/cancers13061291>. <https://www.mdpi.com/2072-6694/13/6/1291>.
- [17] V.M. Campello, P. Gkontra, C. Izquierdo, C. Martín-Isla, A. Sojoudi, P.M. Full, K. Maier-Hein, Y. Zhang, Z. He, J. Ma, M. Parreño, A. Albiol, F. Kong, S. C. Shadden, J.C. Acero, V. Sundaresan, M. Saber, M. Elattar, H. Li, B. Menze, F. Khader, C. Haarbarger, C.M. Scannell, M. Veta, A. Carscadden, K. Punithakumar, X. Liu, S.A. Tsafaris, X. Huang, X. Yang, L. Li, X. Zhuang, D. Viladés, M.L. Descalzo, A. Guala, L. La Mura, M.G. Friedrich, R. Garg, J. Lebel, F. Henriques, M. Karakas, E. Çavuş, S.E. Petersen, S. Escalera, S. Seguí, J. F. Rodríguez-Palomares, K. Lekadir, Multi-centre, multi-vendor and multi-disease cardiac segmentation: the m amp;ms challenge, *IEEE Trans. Med. Imag.* (2021), <https://doi.org/10.1109/TMI.2021.3090082>, 1–1.
- [18] D.V. Carvalho, E.M. Pereira, J.S. Cardoso, Machine learning interpretability: a survey on methods and metrics, *Electronics* 8 (2019) 832.
- [19] A. Chattopadhyay, A. Sarkar, P. Howlader, V.N. Balasubramanian, Grad-cam++: generalized gradient-based visual explanations for deep convolutional networks, in: *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2018, pp. 839–847, <https://doi.org/10.1109/WACV.2018.00097>.
- [20] C. Chen, O. Li, D. Tao, A. Barnett, C. Rudin, J. Su, This looks like that: deep learning for interpretable image recognition, in: *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, 2019, pp. 8928–8939.
- [21] Z. Chen, Y. Bei, C. Rudin, Concept whitening for interpretable image recognition, *Nat. Mach. Intell.* 2 (2020) 772–782, <https://doi.org/10.1038/s42256-020-00265-z>.
- [22] A. Chowdhury, A. Santamaría-Pang, J. Kubricht, P. Tu, Emergent symbolic language based deep medical image classification, in: *2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI)*, 2021, pp. 689–692.
- [23] J.R. Clough, I. Oksuz, E. Puyol-Antón, B. Ruijsink, A.P. King, J.A. Schnabel, Global and local interpretability for cardiac mri classification, in: *Medical Image Computing and Computer Assisted Intervention – MICCAI 2019*, Springer International Publishing, Cham, 2019, pp. 656–664.
- [24] J.P. Cohen, R. Brooks, S. En, E. Zucker, A. Pareek, M. Lungren, A. Chaudhari, Gispplanation via Latent Shift: A Simple Autoencoder Approach to Progressive Exaggeration on Chest X-Rays, 2021. *ArXiv abs/2102.09475*.
- [25] V. Couteaux, O. Nempont, G. Pizaine, I. Bloch, Towards interpretability of segmentation networks by analyzing deepdreams, in: K. Suzuki, M. Reyes, T. Syeda-Mahmood, E. Konukoglu, B. Glocker, R. Wiest, Y. Gur, H. Greenspan, A. Madabhushi (Eds.), *Interpretability of Machine Intelligence in Medical Image*

- Computing and Multimodal Learning for Clinical Decision Support, Springer International Publishing, Cham, 2019, pp. 56–63.
- [26] Y. Dai, G. Wang, Analyzing tongue images using a conceptual alignment deep autoencoder, *IEEE Access* 6 (2018) 5962–5972, <https://doi.org/10.1109/ACCESS.2017.2788849>.
- [27] Y. Dai, G. Wang, K.C. Li, Conceptual alignment deep neural networks, *J. Intell. Fuzzy Syst.* 34 (2018) 1631–1642.
- [28] A.J. DeGrave, J.D. Janizek, S.I. Lee, Ai for radiographic covid-19 detection selects shortcuts over signal, *Nat. Mach. Intell.* 3 (2021) 610–619, <https://doi.org/10.1038/s42256-021-00338-7>.
- [29] Y. Ding, J. Sohn, M. Kawczynski, H. Trivedi, R. Harnish, N. Jenkins, D. Lituev, T. Copeland, M. Aboian, C. Aparici, S. Behr, R. Flavell, S.y. Huang, K. Zalocusky, L. Nardo, Y. Seo, R. Hawkins, M. hernandez pampaloni, D. Hadley, B. Franc, A deep learning model to predict a diagnosis of alzheimer disease by using 18 f-dg pet of the brain, *Radiology* 290 (2018) 180958, <https://doi.org/10.1148/radiol.2018180958>.
- [30] N.K. Dinsdale, M. Jenkinson, A.I. Namburete, Deep learning-based unlearning of dataset bias for mri harmonisation and confound removal, *Neuroimage* 228 (2021) 117689, <https://doi.org/10.1016/j.neuroimage.2020.117689>.
- [31] F. Doshi-Velez, B. Kim, Towards a Rigorous Science of Interpretable Machine Learning, *arXiv: Machine Learning*, 2017.
- [32] J.M. Durán, K. Jongsma, Who is afraid of black box algorithms? on the epistemological and ethical basis of trust in medical ai, *J. Med. Ethics* 47 (2021) 329–335.
- [33] F. Eitel, K. Ritter, Testing the robustness of attribution methods for convolutional neural networks in mri-based alzheimer's disease classification, in: K. Suzuki, M. Reyes, T. Syeda-Mahmood, E. Konukoglu, B. Glocker, R. Wiest, Y. Gur, H. Greenspan, A. Madabhushi (Eds.), *Interpretability of Machine Intelligence in Medical Image Computing and Multimodal Learning for Clinical Decision Support*, Springer International Publishing, Cham, 2019, pp. 3–11.
- [34] F. Eitel, K. Ritter, Testing the robustness of attribution methods for convolutional neural networks in mri-based alzheimer's disease classification, in: K. Suzuki, M. Reyes, T. Syeda-Mahmood, E. Konukoglu, B. Glocker, R. Wiest, Y. Gur, H. Greenspan, A. Madabhushi (Eds.), *Interpretability of Machine Intelligence in Medical Image Computing and Multimodal Learning for Clinical Decision Support*, Springer International Publishing, Cham, 2019, pp. 3–11.
- [35] F. Eitel, E. Soehler, J. Bellmann-Strobl, A.U. Brandt, K. Rupprecht, R.M. Giess, J. Kuchling, S. Assayer, M. Weygandt, J.D. Haynes, M. Scheel, F. Paul, K. Ritter, Uncovering convolutional neural network decisions for diagnosing multiple sclerosis on conventional mri using layer-wise relevance propagation, *Neuroimage: Clin.* 24 (2019) 102003, <https://doi.org/10.1016/j.nicl.2019.102003>.
- [36] D. Erhan, Y. Bengio, A. Courville, P. Vincent, Visualizing Higher-Layer Features of a Deep Network, *Technical Report 1341*, University of Montreal, 2009.
- [37] G. Erion, J.D. Janizek, P. Sturmfels, S.M. Lundberg, S.I. Lee, Improving performance of deep learning models with axiomatic attribution priors and expected gradients, *Nat. Mach. Intell.* 3 (2021) 620–631, <https://doi.org/10.1038/s42256-021-00343-w>.
- [38] K. Faust, Q. Xie, D. Han, K. Goyle, Z. Volynskaya, U. Djuric, P. Diamandis, Visualizing histopathologic deep learning classification and anomaly detection using nonlinear feature space dimensionality reduction, *BMC Bioinf.* 19 (2018).
- [39] S.G. Finlayson, H.W. Chung, I.S. Kohane, A.L. Beam, Adversarial Attacks against Medical Deep Learning Systems, 2019 *arXiv:1804.05296*.
- [40] P. Gamble, R. Jaroensri, H. Wang, F. Tan, M. Moran, T. Brown, I. Flament-Auvigne, E.A. Rakha, M. Toss, D.J. Dabbs, P. Regitnig, N. Olson, J.H. Wren, C. Robinson, G.S. Corrado, L.H. Peng, Y. Liu, C.H. Mermel, D.F. Steiner, P.H. C. Chen, Determining breast cancer biomarker status and associated morphological features using deep learning, *Commun. Med.* 1 (2021) 14, <https://doi.org/10.1038/s43856-021-00013-3>.
- [41] R. Geirhos, P. Rubisch, C. Michaelis, M. Bethge, F.A. Wichmann, W. Brendel, Imagenet-trained cnns are biased towards texture; increasing shape bias improves accuracy and robustness, in: 7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6–9, 2019, 2019.
- [42] A. Ghandeharioun, B. Kim, C.L. Li, B. Jou, B. Eoff, R.W. Picard, Dissect: Disentangled Simultaneous Explanations via Concept Traversals, 2021. *ArXiv abs/2105.15164*.
- [43] A. Ghorbani, A. Abid, J. Zou, Interpretation of neural networks is fragile, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, 33, 2019, pp. 3681–3688, <https://doi.org/10.1609/aaai.v33i01.33013681>.
- [44] A. Ghorbani, D. Ouyang, A. Abid, B. He, J.H. Chen, R.A. Harrington, D.H. Liang, E.A. Ashley, J.Y. Zou, Deep learning interpretation of echocardiograms, *npj Digit. Med.* 3 (2020) 10, <https://doi.org/10.1038/s41746-019-0216-8>.
- [45] A. Ghorbani, J. Wexler, J.Y. Zou, B. Kim, Towards automatic concept-based explanations, in: *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8–14, 2019, Vancouver, BC, Canada*, 2019, pp. 9273–9282.
- [46] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, Y. Bengio, Generative adversarial nets, *Adv. Neural Inf. Process. Syst.* 27 (2014).
- [47] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. C. Courville, Y. Bengio, in: *Generative Adversarial Nets*, NIPS, 2014.
- [48] I.J. Goodfellow, J. Shlens, C. Szegedy, Explaining and Harnessing Adversarial Examples, 2015 *arXiv:1412.6572*.
- [49] M. Graziani, V. Andrearczyk, H. Müller, Regression Concept Vectors for Bidirectional Explanations in Histopathology, 2018. *ArXiv abs/1904.04520*.
- [50] P. Grossmann, O. Stringfield, N. El-Hachem, M.M. Bui, E.R. Velazquez, C. Parmar, R.T.H. Leijenaar, B. Haibe-Kains, P. Lambin, R.J. Gillies, H.J. Aerts, Defining the biological basis of radiomic phenotypes in lung cancer, *eLife* 6 (2017).
- [51] J. He, L. Shang, H. Ji, X. Zhang, Deep learning features for lung adenocarcinoma classification with tissue pathology images, in: D. Liu, S. Xie, Y. Li, D. Zhao, E.S. M. El-Alfy (Eds.), *Neural Information Processing*, Springer International Publishing, Cham, 2017, pp. 742–751.
- [52] I. Higgins, D. Amos, D. Pfau, S. Racaniere, L. Matthey, D. Rezende, A. Lerchner, Towards a Definition of Disentangled Representations, 2018 *arXiv:1812.02230*.
- [53] A. Hinterreiter, M. Streit, B. Kainz, Projective latent interventions for understanding and fine-tuning classifiers, in: J. Cardoso, H. Van Nguyen, N. Heller, P. Henriques Abreu, I. Işgum, W. Silva, R. Cruz, J. Pereira Amorim, V. Patel, B. Roysam, K. Zhou, S. Jiang, N. Le, K. Luu, R. Sznitman, Y. Cheplygina, D. Mateus, E. Trucco, S. Abbasi (Eds.), *Interpretable and Annotation-Efficient Learning for Medical Image Computing*, Springer International Publishing, Cham, 2020, pp. 13–22.
- [54] A. Hoffmann, C. Fanconi, R. Rade, J. Kohler, This Looks like That... does it? Shortcomings of Latent Space Prototype Interpretability in Deep Networks, 2021 *arXiv:2105.02968*.
- [55] A. Holzinger, A.M. Carrington, H. Müller, Measuring the quality of explanations: the system causability scale (scs), *Kunstliche Intelligenz* 34 (2020) 193–198.
- [56] A. Holzinger, B. Malle, A. Saranti, B. Pfeifer, Towards multi-modal causability with graph neural networks enabling information fusion for explainable ai, *Inf. Fusion* 71 (2021) 28–37, <https://doi.org/10.1016/j.inffus.2021.01.008>.
- [57] M. Izadyazdanabadi, E. Belykh, C. Cavallo, X. Zhao, S. Gandhi, L. Moreira, J. Eschbacher, P. Nakaji, M. Preul, Y. Yang, Weakly-supervised learning-based feature localization in confocal laser endomicroscopy glioma images, in: *MICCAI*, 2018.
- [58] A. Janik, J. Dodd, G. Ifrim, K. Sankaran, K. Curran, Interpretability of a deep learning model in the application of cardiac MRI segmentation with an ACDC challenge dataset, in: I. Işgum, B.A. Landman (Eds.), *Medical Imaging 2021: Image Processing*, International Society for Optics and Photonics, SPIE, 2021, pp. 861–872.
- [59] S. Jetley, N.A. Lord, N. Lee, P. Torr, Learn to pay attention, in: *International Conference on Learning Representations*, 2018. <https://openreview.net/forum?id=HyzbhfWRW>.
- [60] W. Jin, X. Li, G. Hamarneh, One Map Does Not Fit All: Evaluating Saliency Map Explanation on Multi-Modal Medical Images, 2021 *arXiv:2107.05047*.
- [61] T. Karras, S. Laine, T. Aila, A style-based generator architecture for generative adversarial networks, in: 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2019, pp. 4396–4405.
- [62] A. Katzmann, O. Taubmann, S. Ahmad, A. Mühlberg, M. Sühling, H.M. Groß, Explaining clinical decision support systems in medical imaging using cycle-consistent activation maximization, *Neurocomputing* 458 (2021) 141–156, <https://doi.org/10.1016/j.neucom.2021.05.081>.
- [63] D. Kaur, S. Uslu, A. Durresti, S. Badve, M. Dunder, Trustworthy explainability acceptance: a new metric to measure the trustworthiness of interpretable ai medical diagnostic systems, in: L. Barolli, K. Yim, T. Enokido (Eds.), *Complex, Intelligent and Software Intensive Systems*, Springer International Publishing, Cham, 2021, pp. 35–46.
- [64] D.S. Kermany, M. Goldbaum, W. Cai, C.C. Valentim, H. Liang, S.L. Baxter, A. McKeown, G. Yang, X. Wu, F. Yan, J. Dong, M.K. Prasadha, J. Pei, M.Y. Ting, J. Zhu, C. Li, S. Hewett, J. Dong, I. Ziyar, A. Shi, R. Zhang, L. Zheng, R. Hou, W. Shi, X. Fu, Y. Duan, V.A. Huu, C. Wen, E.D. Zhang, C.L. Zhang, O. Li, X. Wang, M.A. Singer, X. Sun, J. Xu, A. Tafreshi, M.A. Lewis, H. Xia, K. Zhang, Identifying medical diagnoses and treatable diseases by image-based deep learning, *Cell* 172 (2018) 1122–1131, <https://doi.org/10.1016/j.cell.2018.02.010>, e9.
- [65] B. Kim, M. Wattenberg, J. Gilmer, C.J. Cai, J. Wexler, F.B. Viégas, R. Sayres, Interpretability beyond Feature Attribution: Quantitative Testing with Concept Activation Vectors (Tcav), *ICML*, 2018.
- [66] E. Kim, S. Kim, M. Seo, S. Yoon, Xprotonet: diagnosis in chest radiography with global and local explanations, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021, pp. 15719–15728.
- [67] P.J. Kindermans, K.T. Schütt, M. Alber, K.R. Müller, D. Erhan, B. Kim, S. Dähne, Learning how to explain neural networks: patternnet and patternattribution, in: *International Conference on Learning Representations*, 2018. <https://openreview.net/forum?id=Hkn7CBAwTW>.
- [68] D.P. Kingma, M. Welling, Auto-encoding Variational Bayes, 2014 *arXiv: 1312.6114*.
- [69] P.W. Koh, T. Nguyen, Y.S. Tang, S. Mussmann, E. Pierson, B. Kim, P. Liang, Concept bottleneck models, in: III, H.D., A. Singh (Eds.), *Proceedings of the 37th International Conference on Machine Learning, PMLR*, 2020, pp. 5338–5348, in: <http://proceedings.mlr.press/v119/koh20a.html>.
- [70] A. Krizhevsky, I. Sutskever, G.E. Hinton, Imagenet classification with deep convolutional neural networks, *Adv. Neural Inf. Process. Syst.* 25 (2012) 1097–1105.
- [71] R. LaLonde, D. Torigian, U. Bagci, Encoding visual attributes in capsules for explainable medical diagnoses, in: *Medical Image Computing and Computer Assisted Intervention – MICCAI 2020 - 23rd International Conference*, Proceedings, 2020, pp. 294–304, [https://doi.org/10.1007/978-3-030-59710-8\\_29](https://doi.org/10.1007/978-3-030-59710-8_29).
- [72] P. Lambin, R. Leijenaar, T. Deist, J. Peerlings, E.D. Jong, J.V. Timmeren, S. Sanduleanu, R. Larue, A. Even, A. Jochems, Y.V. Wijk, H. Woodruff, J.V. Soest, T. Lustberg, E. Roelofs, W.V. Elmpot, A. Dekker, F. Mottaghy, J. Wildberger, S. Walsh, Radiomics: the bridge between medical imaging and personalized medicine, *Nat. Rev. Clin. Oncol.* 14 (2017) 749–762.

- [73] P. Lambin, E. Rios-Velazquez, R.T.H. Leijenaar, S. Carvalho, R.G. van Stiphout, P. V. Granton, C.M.L. Zegers, R.J. Gillies, R. Boellard, A. Dekker, H.J. Aerts, Radiomics: extracting more information from medical images using advanced feature analysis, *Eur. J. Cancer* 48 (4) (2012) 441–446.
- [74] P. Lambin, R.V. Stiphout, M. Starmans, E. Rios-Velazquez, G. Nalbantov, H. Aerts, E. Roelofs, W.V. Elmp, P. Boutros, P. Granone, V. Valentini, A. Begg, D. Ruysscher, A. Dekker, Predicting outcomes in radiation oncology—multifactorial decision support systems, *Nat. Rev. Clin. Oncol.* 10 (2013) 27–40.
- [75] Y. LeCun, Y. Bengio, G. Hinton, Deep learning, *Nature* 521 (2015) 436–444, <https://doi.org/10.1038/nature14539>.
- [76] H. Lee, S.T. Kim, Y.M. Ro, Generation of multimodal justification using visual word constraint model for explainable computer-aided diagnosis, in: *Interpretability of Machine Intelligence in Medical Image Computing and Multimodal Learning for Clinical Decision Support*, Springer International Publishing, Cham, 2019, pp. 21–29.
- [77] H. Lee, S. Yune, M. Mansouri, M. Kim, S.H. Tajmir, C.E. Guerrier, S.A. Ebert, S. Pomerantz, J. Romero, S. Kamalian, R. Gonzalez, M. Lev, S. Do, An explainable deep-learning algorithm for the detection of acute intracranial haemorrhage from small datasets, *Nat. Biomed. Eng.* 3 (2018) 173–182.
- [78] D. Lenis, D. Major, M. Wimmer, A. Berg, G. Sluiter, K. Bühler, Domain aware medical image classifier interpretation by counterfactual impact analysis, in: A. L. Martel, P. Abolmaesumi, D. Stoyanov, D. Mateus, M.A. Zuluaga, S.K. Zhou, D. Racoceanu, L. Joskowicz (Eds.), *Medical Image Computing and Computer Assisted Intervention – MICCAI 2020*, Springer International Publishing, Cham, 2020, pp. 315–325.
- [79] K. Li, Z. Wu, K. Peng, J. Ernst, Y. Fu, Tell me where to look: guided attention inference network, in: 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2018, pp. 9215–9223.
- [80] O. Li, H. Liu, C. Chen, C. Rudin, Deep learning for case-based reasoning through prototypes: a neural network that explains its predictions, in: S.A. McIlraith, K. Q. Weinberger (Eds.), *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), the 30th Innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18)*, New Orleans, AAAI Press, Louisiana, USA, 2018, pp. 3530–3537. February 2-7, 2018.
- [81] S. Li, M. Dong, G. Du, X. Mu, Attention dense-u-net for automatic breast mass segmentation in digital mammogram, *IEEE Access* 7 (2019) 59037–59047.
- [82] X. Li, N.C. Dvornek, Y. Zhou, J. Zhuang, P. Ventola, J.S. Duncan, Efficient interpretation of deep learning models using graph structure and cooperative game theory: application to asd biomarker discovery, in: A.C.S. Chung, J.C. Gee, P.A. Yushkevich, S. Bao (Eds.), *Information Processing in Medical Imaging*, Springer International Publishing, Cham, 2019, pp. 718–730.
- [83] X. Li, N.C. Dvornek, J. Zhuang, P. Ventola, J.S. Duncan, Brain biomarker interpretation in asd using deep learning and fmri, in: A.F. Frangi, J.A. Schnabel, C. Davatzikos, C. Alberola-López, G. Fichtinger (Eds.), *Medical Image Computing and Computer Assisted Intervention – MICCAI 2018*, Springer International Publishing, Cham, 2018, pp. 206–214.
- [84] X. Li, J. Wu, E.Z. Chen, H. Jiang, From deep learning towards finding skin lesion biomarkers, in: 2019 41st Annual International Conference of the IEEE Engineering in Medicine and Biology Society, (EMBC), 2019, pp. 2797–2800, <https://doi.org/10.1109/EMBC.2019.8857334>.
- [85] C.C. Liao, Y.F. Chen, F. Xiao, Brain midline shift measurement and its automation: a review of techniques and algorithms, *Int. J. Biomed. Imag.* 2018 (2018) 4303161, <https://doi.org/10.1155/2018/4303161>.
- [86] C.Y. Lin, in: *Rouge: A Package for Automatic Evaluation of Summaries*, 2004. ACL 2004.
- [87] Z.C. Lipton, The Doctor Just Won't Accept that!, 2017 arXiv:1711.08037.
- [88] Z.C. Lipton, The myths of model interpretability, *Queue* 16 (2018) 31–57.
- [89] A. Lopatina, S. Ropele, R. Sibgatulin, J.R. Reichenbach, D. Güllmar, Investigation of deep-learning-driven identification of multiple sclerosis patients based on susceptibility-weighted images using relevance analysis, 609468–609468, *Front. Neurosci.* 14 (2020), <https://pubmed.ncbi.nlm.nih.gov/33390890>.
- [90] S.M. Lundberg, S.I. Lee, A unified approach to interpreting model predictions, in: *Proceedings of the 31st International Conference on Neural Information Processing Systems*, Curran Associates Inc., Red Hook, NY, USA, 2017, pp. 4768–4777.
- [91] G. M. A. V., M.M. S. M. H. Concept attribution: explaining cnn decisions to physicians, *Comput. Biol. Med.* 123 (2020) 103865, <https://doi.org/10.1016/j.combiomed.2020.103865>. <https://www.sciencedirect.com/science/article/pii/S0010482520302225>.
- [92] L. van der Maaten, G. Hinton, Visualizing data using t-sne, *J. Mach. Learn. Res.* 9 (2008) 2579–2605. <http://jmlr.org/papers/v9/vandermaaten08a.html>.
- [93] P.R. Magesh, R.D. Myloth, R.J. Tom, An explainable machine learning model for early detection of Parkinson's disease using lime on datscan imagery, *Comput. Biol. Med.* 126 (2020) 104041, <https://doi.org/10.1016/j.combiomed.2020.104041>.
- [94] A. Mahinpei, J. Clark, I. Lage, F. Doshi-Velez, W. Pan, Promises and Pitfalls of Black-Box Concept Learning Models, 2021 arXiv:2106.13314.
- [95] A. Margelou, M. Ashman, U. Bhatt, Y. Chen, M. Jamnik, A. Weller, Do Concept Bottleneck Models Learn as Intended?, 2021. CoRR abs/2105.04289.
- [96] S. Mohammadjafari, M. Cevik, M. Thanabalasingam, A. Basar, A.D.N. Initiative, Using protopnet for interpretable alzheimer's disease classification, in: *Proceedings of the Canadian Conference on Artificial Intelligence*, 2021, <https://doi.org/10.21428/594757db.fb59ce6c>.
- [97] G. Montavon, S. Lapuschkin, A. Binder, W. Samek, K.R. Müller, Explaining nonlinear classification decisions with deep taylor decomposition, *Pattern Recogn.* 65 (2017) 211–222, <https://doi.org/10.1016/j.patrec.2016.11.008>.
- [98] A. Narayanaswamy, S. Venugopalan, D.R. Webster, L. Peng, G.S. Corrado, P. Ruamviboonsuk, P. Bavishi, M. Brenner, P.C. Nelson, A.V. Varadarajan, Scientific discovery by generating counterfactuals using image translation, in: A. L. Martel, P. Abolmaesumi, D. Stoyanov, D. Mateus, M.A. Zuluaga, S.K. Zhou, D. Racoceanu, L. Joskowicz (Eds.), *Medical Image Computing and Computer Assisted Intervention – MICCAI 2020*, Springer International Publishing, Cham, 2020, pp. 273–283.
- [99] P. Natekar, A. Kori, G. Krishnamurthi, Demystifying brain tumor segmentation networks: interpretability and uncertainty analysis, *Front. Comput. Neurosci.* 14 (2020) 6, <https://doi.org/10.3389/fncom.2020.00006>.
- [100] K.M. Panth, R.T.H. Leijenaar, S. Carvalho, N.G. Liewues, A. Yaromina, L. J. Dubois, P. Lambin, Is there a causal relationship between genetic changes and radiomics-based image features? an in vivo preclinical experiment with doxycycline inducible gadd34 tumor cells, *Radiother. Oncol. : J. Eur. Soc. Therapeut. Radiol. Oncol.* 116 (3) (2015) 462–466.
- [101] H. Panwar, P. Gupta, M.K. Siddiqui, R. Morales-Menendez, P. Bhardwaj, V. Singh, A deep learning and grad-cam based color visualization approach for fast detection of covid-19 cases using chest x-ray and ct-scan images, *Chaos, Solit. Fractals* 140 (2020) 110190, <https://doi.org/10.1016/j.chaos.2020.110190>.
- [102] Z. Papanastasiopoulos, R.K. Samala, H.P. Chan, L. Hadjiiski, C. Paramagul, M. A. Helvie, C.H. Neal, Explainable AI for medical imaging: deep-learning CNN ensemble for classification of estrogen receptor status from breast MRI, in: H. K. Hahn, M.A. Mazurowski (Eds.), *Medical Imaging 2020: Computer-Aided Diagnosis*, International Society for Optics and Photonics, SPIE, 2020, pp. 228–235, <https://doi.org/10.1117/12.2549298>.
- [103] K. Papineni, S. Roukos, T. Ward, W.J. Zhu, in: *Bleu: a Method for Automatic Evaluation of Machine Translation*, ACL, 2002.
- [104] S. Pereira, R. Meier, V. Alves, M. Reyes, C.A. Silva, Automatic brain tumor grading from mri data using convolutional neural networks and quality assessment, in: *Understanding and Interpreting Machine Learning in Medical Image Computing Applications*, Springer International Publishing, Cham, 2018, pp. 106–114.
- [105] M. Pisov, M. Goncharov, N. Kurochkina, S. Morozov, V. Gombolevskiy, V. Chernina, A. Vladzimirskiy, K. Zamyatina, A. Chesnokova, I. Pronin, M. Shifrin, M. Belyaev, Incorporating task-specific structural knowledge into cnns for brain midline shift detection, in: *Interpretability of Machine Intelligence in Medical Image Computing and Multimodal Learning for Clinical Decision Support*, Springer International Publishing, Cham, 2019, pp. 30–38.
- [106] E. Puyol-Antón, C. Chen, J.R. Clough, B. Ruijsink, B.S. Sidhu, J. Gould, B. Porter, M. Elliott, V. Mehta, D. Rueckert, C.A. Rinaldi, A.P. King, Interpretable deep models for cardiac resynchronisation therapy response prediction, in: *Medical Image Computing and Computer Assisted Intervention – MICCAI 2020*, Springer International Publishing, 2020, pp. 284–293.
- [107] S. Ravi, S. Khoshrou, M. Pechenizkiy, Vidi: Descriptive Visual Data Clustering as Radiologist Assistant in Covid-19 Streamline Diagnostic, 2020. ArXiv abs/2011.14871.
- [108] M.T. Ribeiro, S. Singh, C. Guestrin, in: "why Should I Trust You?": Explaining the Predictions of Any Classifier. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2016.
- [109] L. Rieger, C. Singh, W.J. Murdoch, B. Yu, in: *Interpretations Are Useful: Penalizing Explanations to Align Neural Networks with Prior Knowledge*, ICML, 2020.
- [110] C. Rudin, Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead, *Nat. Mach. Intell.* 1 (2019) 206–215, <https://doi.org/10.1038/s42256-019-0048-x>.
- [111] S. Sabour, N. Frosst, G.E. Hinton, Dynamic routing between capsules, in: I. Guyon, U.V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, R. Garnett (Eds.), *Advances in Neural Information Processing Systems*, Curran Associates, Inc, 2017.
- [112] A. Saha, L. Grimm, M.R. Harowicz, S. Ghate, C.E. Kim, R. Walsh, M. Mazurowski, Interobserver variability in identification of breast tumors in mri and its implications for prognostic biomarkers and radiogenomics, *Med. Phys.* 43 8 (2016) 4558.
- [113] Z. Salahuddin, M. Lenga, H. Nickisch, Multi-resolution 3d convolutional neural networks for automatic coronary centerline extraction in cardiac ct angiography scans, in: 2021 IEEE 18th International Symposium on Biomedical Imaging, (ISBI), 2021, pp. 91–95, <https://doi.org/10.1109/ISBI48211.2021.9434002>.
- [114] H. Saleem, A.R. Shahid, B. Raza, Visual interpretability in 3d brain tumor segmentation network, *Comput. Biol. Med.* 133 (2021) 104410, <https://doi.org/10.1016/j.combiomed.2021.104410>.
- [115] W. Samek, A. Binder, G. Montavon, S. Lapuschkin, K.R. Müller, Evaluating the visualization of what a deep neural network has learned, *IEEE Trans. Neural Networks Learn. Syst.* 28 (2017) 2660–2673, <https://doi.org/10.1109/TNNLS.2016.2599820>.
- [116] S. Sanduleanu, H.C. Woodruff, E.E. de Jong, J.E. van Timmeren, A. Jochems, L. Dubois, P. Lambin, Tracking tumor biology with radiomics: a systematic review utilizing a radiomics quality score, *Radiother. Oncol.* 127 (2018) 349–360, <https://doi.org/10.1016/j.radonc.2018.03.033>.
- [117] A. Santamaría-Pang, J. Kubricht, A. Chowdhury, C. Bhushan, P.H. Tu, Towards emergent language symbolic semantic segmentation and model interpretability, in: *Medical Image Computing and Computer Assisted Intervention – MICCAI 2020 - 23rd International Conference*, Lima, Peru, October 4-8, 2020, Proceedings, Part 1, Springer, 2020, pp. 326–334, [https://doi.org/10.1007/978-3-030-59710-8\\_32](https://doi.org/10.1007/978-3-030-59710-8_32).

- [118] R. Sayres, A. Taly, E. Rahimy, K. Blumer, D. Coz, N. Hammel, J. Krause, A. Narayanaswamy, Z. Rastegar, D. Wu, S. Xu, S. Barb, A. Joseph, M. Shumski, J. Smith, A.B. Sood, G.S. Corrado, L. Peng, D.R. Webster, Using a deep learning algorithm and integrated gradients explanation to assist grading for diabetic retinopathy, *Ophthalmology* 126 (2019) 552–564, <https://doi.org/10.1016/j.optha.2018.11.016>.
- [119] J. Schlemper, O. Oktay, M. Schaap, M. Heinrich, B. Kainz, B. Glocker, D. Rueckert, Attention gated networks: learning to leverage salient regions in medical images, *Med. Image Anal.* 53 (2019) 197–207.
- [120] J. Schmidhuber, Deep learning in neural networks: an overview, *Neural Network* 61 (2015) 85–117, <https://doi.org/10.1016/j.neunet.2014.09.003>. <https://www.sciencedirect.com/science/article/pii/S0893608014002135>.
- [121] K. Schutte, O. Moindrot, P. Hérent, J.B. Schiratti, S. Jégou, Using Stylegan for Visual Interpretability of Deep Learning Models on Medical Images, 2021. ArXiv abs/2101.07563.
- [122] J. Seah, J.S.N. Tang, A. Kitchen, F. Gaillard, A.F. Dixon, Chest radiographs in congestive heart failure: visualizing neural network learning, *Radiology* 290 (2) (2019) 514–522.
- [123] P. Seegerer, A. Binder, R. Saitenmacher, M. Bockmayr, M. Alber, P. Jurmeister, F. Klauschen, K.R. Müller, Interpretable Deep Neural Network to Predict Estrogen Receptor Status from Haematoxylin-Eosin Images, Springer International Publishing, Cham, 2020, pp. 16–37, [https://doi.org/10.1007/978-3-030-50402-1\\_2](https://doi.org/10.1007/978-3-030-50402-1_2).
- [124] R.R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, D. Batra, Grad-cam: Visual Explanations from Deep Networks via Gradient-Based Localization, in: 2017 IEEE International Conference on Computer Vision (ICCV), 2017, pp. 618–626, <https://doi.org/10.1109/ICCV.2017.74>.
- [125] S. Shen, S.X. Han, D.R. Aberler, A.A. Bui, W. Hsu, An interpretable deep hierarchical semantic convolutional neural network for lung nodule malignancy classification, *Expert Syst. Appl.* 128 (2019) 84–95, <https://doi.org/10.1016/j.eswa.2019.01.048>.
- [126] A. Shrikumar, P. Greenside, A. Shcherbina, A. Kundaje, Not Just a Black Box: Learning Important Features through Propagating Activation Differences, 2017 arXiv:1605.01713.
- [127] W. Silva, A. Poellinger, J.S. Cardoso, M. Reyes, Interpretability-guided content-based medical image retrieval, in: *Medical Image Computing and Computer Assisted Intervention – MICCAI 2020*, Springer International Publishing, Cham, 2020, pp. 305–314.
- [128] K. Simonyan, A. Vedaldi, A. Zisserman, Deep inside convolutional networks: visualising image classification models and saliency maps, in: 2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14–16, 2014, Workshop Track Proceedings, 2014.
- [129] G. Singh, K. Yow, These do not look like those: an interpretable deep learning model for image recognition, *IEEE Access* 9 (2021) 41482–41493.
- [130] G. Singh, K.C. Yow, An interpretable deep learning model for covid-19 detection with chest x-ray images, *IEEE Access* 9 (2021) 85198–85208, <https://doi.org/10.1109/ACCESS.2021.3087583>.
- [131] S. Singla, B. Pollack, S. Wallace, K. Batmanghelich, Explaining the Black-Box Smoothly- a Counterfactual Approach, 2021 arXiv:2101.04230.
- [132] D. Smilkov, N. Thorat, B. Kim, F. Viégas, M. Wattenberg, Smoothgrad: Removing Noise by Adding Noise, 2017 arXiv:1706.03825.
- [133] R. Smith-Bindman, D.L. Miglioretti, E.B. Larson, Rising use of diagnostic medical imaging in a large integrated health system, *Health Aff.(Project Hope)* 27 (2008) 1491–1502, <https://doi.org/10.1377/hlthaff.27.6.1491>. <https://pubmed.ncbi.nlm.nih.gov/18997204>.
- [134] J.T. Springenberg, A. Dosovitskiy, T. Brox, M.A. Riedmiller, Striving for Simplicity: the All Convolutional Net, 2015. CoRR abs/1412.6806.
- [135] A. Stergiou, G. Kapidis, G. Kalliatakis, C. Chrysoulas, R. Veltkamp, R. Poppe, Saliency tubes: visual explanations for spatio-temporal convolutions, in: 2019 IEEE International Conference on Image Processing (ICIP), IEEE, 2019, pp. 1830–1834.
- [136] J. Sun, F. Darbehani, M. Zaidi, B. Wang, Saunet: shape attentive u-net for interpretable medical image segmentation, in: *Medical Image Computing and Computer Assisted Intervention – MICCAI 2020*, Springer International Publishing, Cham, 2020, pp. 797–806.
- [137] M. Sundararajan, A. Taly, Q. Yan, Axiomatic attribution for deep networks, in: *Proceedings of the 34th International Conference on Machine Learning*, vol. 70, 2017, pp. 3319–3328.
- [138] Y. Tang, Y. Tang, Y. Zhu, J. Xiao, R.M. Summers, A disentangled generative model for disease decomposition in chest x-rays via normal image synthesis, *Med. Image Anal.* 67 (2021) 101839, <https://doi.org/10.1016/j.media.2020.101839>.
- [139] Z. Tang, K. Chuang, C. DeCarli, L.W. Jin, L. Beckett, M. Keiser, B. Dugger, Interpretable classification of alzheimer's disease pathologies with a convolutional neural network pipeline, *Nat. Commun.* 10 (2019), <https://doi.org/10.1038/s41467-019-10212-1>.
- [140] M. Temme, Algorithms and transparency in view of the new general data protection regulation, *Eur. Data Prot. L. Rev.* 3 (2017) 473.
- [141] A.W. Thomas, H. Heekeren, K. Müller, W. Samek, Analyzing neuroimaging data through recurrent deep learning models, *Front. Neurosci.* 13 (2019).
- [142] E. Tjoa, C. Guan, Quantifying Explainability of Saliency Methods in Deep Neural Networks, 2020. ArXiv abs/2009.02899.
- [143] P. Tschandl, N. Codella, B.N. Akay, G. Argenziano, R. Braun, H. Cabo, D. Gutman, A. Halpern, B. Helba, R. Hofmann-Wellenhof, A. Lallas, J. Lapins, C. Longo, J. Malvehy, M. Marchetti, A. Marghoob, S. Menzies, A. Oakley, J. Paoli, S. Puig, C. Rinner, C. Rosendahl, A. Scope, C. Sinz, H. Soyer, L. Thomas, I. Zalaudek, H. Kittler, Comparison of the accuracy of human readers versus machine-learning algorithms for pigmented skin lesion classification: an open, web-based, international, diagnostic study, *Lancet Oncol.* 20 (7) (2019) 938–947.
- [144] P. Tschandl, C. Rinner, Z. Apalla, G. Argenziano, N. Codella, A. Halpern, M. Janda, A. Lallas, C. Longo, J. Malvehy, J. Paoli, S. Puig, C. Rosendahl, H. Soyer, I. Zalaudek, H. Kittler, Human-computer collaboration for skin cancer recognition, *Nat. Med.* (2020) 1–6.
- [145] H. Uzunova, J. Ehrhardt, T. Kepp, H. Handels, Interpretable explanations of black box classifiers applied on medical images by meaningful perturbations using variational autoencoders, in: E.D. Angelini, B.A. Landman (Eds.), *Medical Imaging 2019: Image Processing*, International Society for Optics and Photonics, SPIE, 2019, pp. 264–271, <https://doi.org/10.1117/12.2511964>.
- [146] R. Vedantam, C.L. Zitnick, D. Parikh, Cider: consensus-based image description evaluation, in: IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2015, IEEE Computer Society, Boston, MA, USA, 2015, pp. 4566–4575. June 7–12, 2015.
- [147] S. Verma, J.P. Dickerson, K.E. Hines, Counterfactual Explanations for Machine Learning: A Review, 2020. ArXiv abs/2010.10596.
- [148] S. Waldstein, P. Seeböck, R. Donner, A. Sadeghipour, H. Bogunović, A. Osborne, U. Schmidt-Erfurth, Unbiased identification of novel subclinical imaging biomarkers using unsupervised deep learning, *Sci. Rep.* 10 (2020) 12954, <https://doi.org/10.1038/s41598-020-69814-1>.
- [149] X. Wang, Y. Peng, L. Lu, Z. Lu, R. Summers, Tienet: text-image embedding network for common thorax disease classification and reporting in chest x-rays, in: 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2018, pp. 9049–9058.
- [150] V. Wargnier-Dauchelle, T. Grenier, F. Durand-Dubief, F. Cotton, M. Sdika, A more interpretable classifier for multiple sclerosis, in: 2021 IEEE 18th International Symposium on Biomedical Imaging, ISBI, 2021, pp. 1062–1066, <https://doi.org/10.1109/ISBI48211.2021.9434074>.
- [151] O. Weaver, J. Leung, Biomarkers and imaging of breast cancer, *Am. J. Roentgenol.* 210 (2017) 1–8, <https://doi.org/10.2214/AJR.17.18708>.
- [152] K. Wickstrom, M. Kampffmeyer, R. Jensen, Uncertainty and interpretability in convolutional neural networks for semantic segmentation of colorectal polyps, *Med. Image Anal.* 60 (2020) 101619, <https://doi.org/10.1016/j.media.2019.101619>. <https://www.sciencedirect.com/science/article/pii/S1361841519301574>.
- [153] W. Yan, S. Plis, V.D. Calhoun, S. Liu, R. Jiang, T.Z. Jiang, J. Sui, Discriminating schizophrenia from normal controls using resting state functional network connectivity: a deep neural network and layer-wise relevance propagation method, in: 2017 IEEE 27th International Workshop on Machine Learning for Signal Processing, MLSP, 2017, pp. 1–6.
- [154] Y. Yan, J. Kawahara, G. Hamarneh, Melanoma recognition via visual attention, in: A.C.S. Chung, J.C. Gee, P.A. Yushkevich, S. Bao (Eds.), *Information Processing in Medical Imaging*, Springer International Publishing, Cham, 2019, pp. 793–804.
- [155] G. Yang, Q. Ye, J. Xia, Unbox the Black-Box for the Medical Explainable AI via Multi-Modal and Multi-Centre Data Fusion: A Mini-Review, Two Showcases and beyond, 2021. CoRR abs/2102.01998.
- [156] H. Yang, J.Y. Kim, H. Kim, S. Adhikari, Guided soft attention network for classification of breast cancer histopathology images, *IEEE Trans. Med. Imag.* 39 (2020) 1306–1315.
- [157] J. Yang, N. Dvornek, F. Zhang, J. Zhuang, J. Chapiro, M. Lin, J. Duncan, Domain-agnostic learning with anatomy-consistent embedding for cross-modality liver segmentation, in: 2019 IEEE/CVF International Conference on Computer Vision Workshop (ICCVW), 2019, pp. 323–331.
- [158] J. Yang, H. Veeraraghavan, S. Armato, K. Farahani, J. Kirby, J. Kalpathy-Kramer, W.V. van Elmpt, A. Dekker, X. Han, X. Peng, P. Aljabar, B. Oliveira, B. van der Heyden, L. Zamdborg, D. Lam, M. Gooding, G. Sharp, Autosegmentation for thoracic radiation treatment planning: a grand challenge at aapm 2017, *Med. Phys.* 45 (2018) 4568–4581.
- [159] H. Yeche, J. Harrison, T. Berthier, Ubs: a dimension-agnostic metric for concept vector interpretability applied to radiomics, in: *iMIMIC/ML-CDS@MICCAI*, 2019.
- [160] E.B. Yilmaz, A.O. Mader, T. Fricke, J. Peña, C.C. Glüer, C. Meyer, Assessing attribution maps for explaining cnn-based vertebral fracture classifiers, in: *Interpretable and Annotation-Efficient Learning for Medical Image Computing*, Springer International Publishing, Cham, 2020, pp. 3–12.
- [161] K. Young, G. Booth, B. Simpson, R. Dutton, S. Shrapnel, Deep neural network or dermatologist? in: K. Suzuki, M. Reyes, T. Syeda-Mahmood, E. Konukoglu, B. Glocker, R. Wiest, Y. Gur, H. Greenspan, A. Madabhushi (Eds.), *Interpretability of Machine Intelligence in Medical Image Computing and Multimodal Learning for Clinical Decision Support* Springer International Publishing, Cham, 2019, pp. 48–55.
- [162] M.D. Zeiler, R. Fergus, Visualizing and understanding convolutional networks, in: D. Fleet, T. Pajdla, B. Schiele, T. Tuytelaars (Eds.), *Computer Vision – ECCV 2014*, Springer International Publishing, Cham, 2014, pp. 818–833.
- [163] J. Zhang, C. Petitjean, F. Yger, S. Aïnouz, Explainability for regression cnn in fetal head circumference estimation from ultrasound images, in: *Interpretable and Annotation-Efficient Learning for Medical Image Computing*, Springer International Publishing, Cham, 2020, pp. 73–82.
- [164] R. Zhang, S. Tan, R. Wang, S. Manivannan, J. Chen, H. Lin, W.S. Zheng, Biomarker localization by combining cnn classifier and generative adversarial network, in: D. Shen, T. Liu, T.M. Peters, L.H. Staib, C. Essert, S. Zhou, P.T. Yap, A. Khan (Eds.), *Medical Image Computing and Computer Assisted Intervention – MICCAI 2019*, Springer International Publishing, Cham, 2019, pp. 209–217.
- [165] Z. Zhang, P. Chen, M. McGough, F. Xing, C. Wang, M. Bui, Y. Xie, M. Sapkota, L. Cui, J. Dhillon, N. Ahmad, F. Khalil, S. Dickinson, X. Shi, F. Liu, H. Su, J. Cai,

- L. Yang, Pathologist-level interpretable whole-slide cancer diagnosis with deep learning, *Nat. Mach. Intell.* 1 (2019) 236–245.
- [166] Z. Zhang, Y. Xie, F. Xing, M. McGough, L. Yang, Mdnet: a semantically and visually interpretable medical image diagnosis network, in: 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2017, pp. 3549–3557.
- [167] G. Zhao, B. Zhou, K. Wang, R. Jiang, M. Xu, in: Respond-cam: Analyzing Deep Models for 3d Imaging Data by Visualizations, MICCAI, 2018.
- [168] B. Zhou, A. Khosla, A. Lapedriza, A. Oliva, A. Torralba, Learning deep features for discriminative localization, in: 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2016, pp. 2921–2929.
- [169] J. Zhou, A.H. Gandomi, F. Chen, A. Holzinger, Evaluating the quality of machine learning explanations: a survey on methods and metrics, *Electronics* 10 (2021).
- [170] J.Y. Zhu, T. Park, P. Isola, A.A. Efros, Unpaired image-to-image translation using cycle-consistent adversarial networks, in: 2017 IEEE International Conference on Computer Vision (ICCV), 2017, pp. 2242–2251, <https://doi.org/10.1109/ICCV.2017.244>.
- [171] P. Zhu, M. Ogino, Guideline-based additive explanation for computer-aided diagnosis of lung nodules, in: iMIMIC/ML-CDS@MICCAI, 2019.
- [172] J. Zhuang, N.C. Dvornek, X. Li, P. Ventola, J.S. Duncan, Invertible network for classification and biomarker selection for ASD, in: D. Shen, T. Liu, T.M. Peters, L. H. Staib, C. Essert, S. Zhou, P.T. Yap, A. Khan (Eds.), *Medical Image Computing and Computer Assisted Intervention – MICCAI 2019*, Springer International Publishing, Cham, 2019, pp. 700–708.
- [173] X. Zhuang, L. Li, C. Payer, D. Stern, M. Urschler, M.P. Heinrich, J. Oster, C. Wang, Örjan Smedby, C. Bian, X. Yang, P.A. Heng, A. Mortazi, U. Bagci, G. Yang, C. Sun, G. Galisot, J.Y. Ramel, T. Brouard, Q. Tong, W. Si, X. Liao, G. Zeng, Z. Shi, G. Zheng, C. Wang, T. MacGillivray, D. Newby, K. Rhode, S. Ourselin, R. Mohiaddin, J. Keegan, D. Firmin, G. Yang, Evaluation of algorithms for multi-modality whole heart segmentation: an open-access grand challenge, *Med. Image Anal.* 58 (2019) 101537, <https://doi.org/10.1016/j.media.2019.101537>.