

Interpreting Deep Machine Learning Models: An Easy Guide for Oncologists

José P. Amorim, Pedro H. Abreu, Alberto Fernández, Mauricio Reyes, João Santos, Miguel H. Abreu

Abstract—Healthcare agents, in particular in the oncology field, are currently collecting vast amounts of diverse patient data. In this context, some decision-support systems, mostly based on deep learning techniques, have already been approved for clinical purposes. Despite all the efforts in introducing artificial intelligence methods in the workflow of clinicians, its lack of interpretability - understand how the methods make decisions - still inhibits their dissemination in clinical practice. The aim of this article is to present an easy guide for oncologists explaining how these methods make decisions and illustrating the strategies to explain them. Theoretical concepts were illustrated based on oncological examples and a literature review of research works was performed from PubMed between January 2014 to September 2020, using “deep learning techniques”, “interpretability” and “oncology” as keywords. Overall, more than 60% are related to breast, skin or brain cancers and the majority focused on explaining the importance of tumor characteristics (e.g. dimension, shape) in the predictions. The most used computational methods are multilayer perceptrons and convolutional neural networks. Nevertheless, despite being successfully applied in different cancers scenarios, endowing deep learning techniques with interpretability, while maintaining their performance, continues to be one of the greatest challenges of artificial intelligence.

Index Terms—big data, interpretability, deep learning, decision-support systems, oncology

I. INTRODUCTION

Today, in healthcare scenarios, we are living in a digital era where physical patient records are mapped to digital formats. This has opened the possibility to improve the efficiency and quality of treatment provided to patients by building decision-support systems.

Machine Learning (ML) is a sub-field of Artificial Intelligence (AI) which studies algorithms that are capable to construct data driven models. The construction of such models follows two distinct phases - training and application. During training, the algorithm builds a model which fits the data received as input, while in application, the now trained model will produce results based on a new set of information that

it receives exclusively in this phase and can be used to test its performance.

Between 2014 and 2019 the US Food and Drug Administration approved 46 ML algorithms [1] for clinical purposes encompassing different areas like mammogram screening and ultrasound image diagnosis, turning the application of ML in healthcare context a reality.

The majority of these algorithms are supervised which means that in these scenarios, they need a help of a physician to label the data before the mining process starts. As an example, in overall survival prediction of breast cancer patients it is necessary that a physician labels the set of patient data that will be used in the training process with the target variable. When this target variable is discrete we are present to a classification problem (benign or malignant), or a regression problem in case the variable is continuous (overall survival - measured in months).

Among different ML paradigms that are used in medical contexts, the Artificial Neural Network (ANN) is a popular supervised algorithm inspired by biological neuron, and began to be used in healthcare in the early 90s [2]. The ANN is an analogy used by computer scientists to emulate the behaviour of the human brain and are composed by an input, an output and intermediate layers, which are also called hidden layers. Similarly to biological neurons, each artificial neuron, or perceptron [3], receives a set of inputs, either from the input layer or other neurons, performs a linear combination based on its weights and make a non-linear decision whether to activate the neuron and fires it.

Due to the increasing computational power, the complexity of these networks has substantially grown, materializing in the use of dozens of layers and millions of neurons. In this context, Deep Learning (DL) techniques - a subset of ANN techniques - emerged as the state of the art for many real world problems, surpassing other ML techniques, and reaching human-level performance in several tasks such as in the classification of melanoma from dermoscopic images [4], or the detection of lymph node metastases in breast cancers from pathology images [5].

Despite its vast potential DL suffers from several disadvantages. First is the dependency on large amounts of data and computational power. Also the black-box nature of DL makes it difficult to interpret their decisions and prevents their dissemination in clinical practice.

The objective of this study is to present an easy guide for oncologists explaining how DL techniques make decisions and illustrating the strategies that can be used in the oncological field to explain them, as it is an essential step towards the

Manuscript submitted March 5, 2021. This work was supported in part by the FCT Research Grant SFRH/BD/136786/2018.

José P. Amorim and Pedro H. Abreu are with the CISUC, Department of Informatics Engineering, University of Coimbra, Coimbra 3030-790, (e-mail: jпаморим@dei.uc.pt).

Alberto Fernández is with the Department of Computer Science and Artificial Intelligence, University of Granada, Granada 18071, Spain.

Mauricio Reyes is with the Insel Data Science Center, Inselspital, Bern University Hospital, Bern 3010, Switzerland and the ARTORG Center for Biomedical Research, University of Bern, Bern 3008, Switzerland.

José P. Amorim, João Santos and Miguel H. Abreu are with the IPO-Porto Research Center - Portuguese Institute of Oncology of Porto, Porto 4200-072, Portugal.

integration of DL in the workflow of physicians in the field of oncology. To better illustrate these strategies to oncologists and other healthcare agents we give self-explanatory oncological examples. Other reviews already covered specific medical areas such as radiology [6] which only equate to a small set of image modalities and does not cover other patient data such as genomic data. Others expand the review to the medical field but do not focus on DL techniques [7]. This is the first study to review in detail work of interpretability of DL techniques in the oncological field.

Results from this study were compiled by searching the PubMed database for articles published between January 2014 and September 2020, searching individually and in combination search terms such as “interpretability”, “deep learning”, “oncology”, “cancer” and “decision support systems”.

Overall, from this selection, more than 60% are related to breast, skin or brain cancers and the majority focused on explaining the importance of tumor characteristics (e.g. dimension, shape) in the disease behavior prediction. Among the DL techniques used in the oncology field which were interpreted, the majority are multilayer perceptrons and convolutional neural networks. In this study we also have found that the majority of works focus on medical imaging (e.g. mammogram, histological images and dermoscopic images) related to breast and skin cancer. Possible explanations are related to the most prevalent diseases and also the dissemination of well curated datasets and challenges target at those diseases. Overall, most works focus on the validation of the knowledge acquired by the DL model for the diagnosis of malignancy or detection of a cancer disease.

Despite being successfully applied in different cancer scenarios, endowing deep learning techniques with the ability to explain their predictions, while maintaining their exceptional performance, will continue to be one of the greatest challenges faced by artificial intelligence. Future work includes the extension of interpretability methods for debugging model misbehavior and acquire new knowledge about disease, as well as largely overlooked cancer tasks such as tumor segmentation and image registration. Also, the evaluation of interpretability methods so that they can be compared and validated.

Throughout the next two overview sections, we will talk about various ANN techniques illustrating their internal architectures and learning processes using a self-explanatory oncological example, that consists of the classification of a breast tumor based on handcrafted features such as mass density (fat-containing - 0, low - 1, equal - 2, high - 3), shape (round - 0, oval - 1, irregular - 2) and the breast side that it was found (left - 0 or right - 1) as well as the raw mammogram. Using such features as an input, the goal of the different types of ANN’s will be predict an output related to the malignancy of the tumor (benign - 0 or malignant - 1).

II. ANN TECHNIQUES OVERVIEW

Artificial Neural Networks (ANN) are a set of algorithms, inspired by the human brain, which are used to approximate unknown functions. They are sometimes called “universal approximators”, because they can learn to approximate mappings

between any input x and any output y , assuming they are correlated. ANNs are composed of layers of neurons, which combine input from the data with a set of coefficients, or weights, assigning significance to inputs with regard to the output label.

Perceptron - The Perceptron [8] is the the precursor to the ANN techniques. In this binary classification algorithm, the linear predictor chooses to “fire” based function combining a set of weights with the input vector.

Training process - As seen in Figure 1, after receiving a set of variables as input (x_1, x_2, \dots, x_n), the perceptron will attribute weights for each variable (w_1, w_2, \dots, w_n) and afterwards will use a mathematical function also known as activation function that will use the weighted sum of the input variables to produce a desired output (y). For each set of input variables, the output (y) is compared to the label corresponding to expected output, also known as target. During training, the weights are continuously changed to move the output of the perceptron and the target closer together.

In the example provided in Figure 1, the perceptron is given the breast cancer tumor variables density, shape, and side and given the weights obtained during training (0.8, 0.7 and 0 respectively), predicts the tumor to be malignant.




Fig. 1: The Perceptron computes the weighted sum of the breast cancer tumor input variables, and an activation function turns the output into a binary prediction of malignancy.

Multilayer Perceptron - The Multilayer Perceptron (MLP) [8] is the natural extension of the perceptron to solve more complex problems. Rather than having a single unit, or neuron, the MLP has multiple layers with multiple neurons each, as can be seen in Figure 2. Also, the linear activation function of the perceptron is replaced by a non-linear activation function which help to solve non-linear problems. Due to its multiple layered structure, the MLP can be seen as a deep neural network.

Training process - After receiving a set of variables as input (x_0, x_1, \dots, x_n), each intermediate neurons present in the hidden layers acts like a perceptron, performing the weighted combination of its inputs and apply a non-linear activation function. The output of activation function of each neuron, also known as activation, acts as input for the neurons of the next layer. The combination of activation of the last intermediate layer produces a desired output (y).

MLPs have been explored on multiple public datasets for breast cancer diagnosis based on tumor characteristics such as density, shape, and side with high accuracy (>97%). Figure 2 illustrates the approach used in [9] based on the public Wisconsin Breast Cancer dataset. In the example, given the tumor

variables (density, shape, and side) the model learns to optimal weight's values during training, to predict the malignancy.




Fig. 2: A Multilayer Perceptron (MLP) composed by an input layer, an output layer and two hidden layers similar to perceptrons that predict malignancy based on breast cancer tumor variables.

Due to their nature, MLPs do not scale well to images. As an example, for an image with a width and height of 100 pixels, the MLP would require 10,000 neurons just in the first layer and this number would grow exponentially with each layer.

Convolutional Neural Networks - Convolutional Neural Networks (CNN) [10], [11] techniques emerged as a solution to addresses the previous computational problem.

Training process - CNNs treat the image as a matrix (Figure 3), extracting features using a mathematical operation called convolution which helps preserve the spatial relationship between neighboring pixels. The convolution slides a small matrix, called a filter, over the original image, and for every position, it computes the element-wise multiplication between the two matrices, and the resulting value forms a single element of the output matrix, called feature map. The filter is composed of weights (w) that are learned during training.

During feature extraction, each convolutional layer is composed by n filters resulting in n feature maps. The values of the feature maps of the last convolutional layer are concatenated into a single vector and used as an input for a MLP which makes the prediction y . During training, the values of the filter matrices and of the MLP are continuously changed to move the output closer to the expected targets.




Fig. 3: The convolution operation produces a feature map, where each element is the result of the element-wise multiplication between the region of the image and the filter (shown in the same shade).




Fig. 4: Representation of Convolutional Neural Network (CNN) used in [5] for the detection of lymph node metastases of breast cancer in histopathological images. First, each convolutional layer produces features maps using the convolution operator across the previous layers' output. The output of the feature extraction is concatenated into a feature vector which serves as input for the classification MLP which predicts the presence of metastases.

CNNs were used for example in the context of detection of lymph node metastases of breast cancer based on whole-slide images of digitally scanned tissue sections of over two hundred patients [5]. Figure 4 illustrates the approach which led to a performance comparable with an expert pathologist interpreting the slides. The CNN learns the weights of the filters, and during the feature extraction is able to extract features which may include the color and shape of the nuclei. The features are used to make the classification, which predicts the tissue to be malignant.

Although CNNs are able to take advantage of the spatial relationships between pixels, they struggle with large sequence data such as text.

Recurrent Neural Networks - Recurrent Neural Networks (RNN) techniques solve this issue by having a small network looped for each element of the sequence, allowing information to persist. A simple RNN contains an hidden state, h_t , at time t which depends on the input of the current step t and the state of the previous step.

Training process - RNNs are usually composed of only a layer of neurons, which taking an input (x_i) predicts the output (o_i) in a recurrent way (Figure 5a).

This refers to the fact that its processing unit (P) is looped n times, where n represents the number of elements of the sequence. During training, the weights of the RNN are continuously changed to minimize the difference between the target sequenced, and the predicted one. As represented by the self-arrow in Figure 5b, the processing unit shares information among steps allowing the context and information from each slice to be passed on until a final diagnosis is given (y_t) [12].

In the example provided in Figure 5a, the RNN is presented in an unfolded version, where the processing unit is repeated for each step in the sequence. It corresponds to an approach for the treatment prognosis of patients with lung cancer based on Computerized Tomography (CT) of four different stages (pre-treatment, 1 month follow-up, 3 month follow-up and 6 months follow-up) [13]. Outcomes such as survival and metastases were predicted using a RNN based on the a set of features extracted from the CT using a CNN. At each step,

and based on the context that is passed from the previous step, it learned to extract and memorize useful context and pass it to subsequent steps until a final prognosis was made.

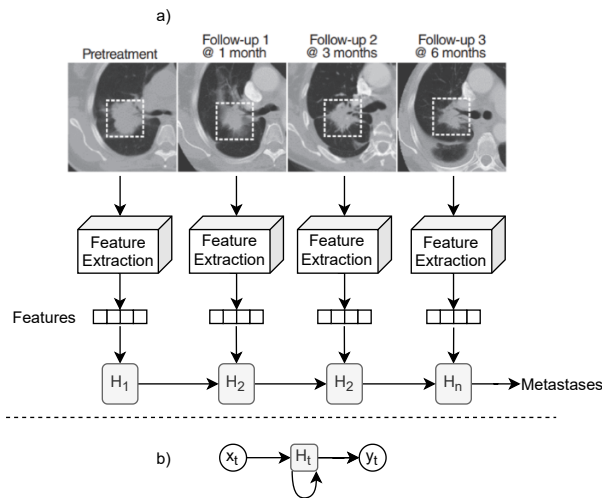


Fig. 5: a) The Recurrent Neural Network (RNN) first extracts set of visual features from CT slides from multiple stages using a CNN. The hidden units optimize their weights to learn useful information from the features and pass stage-specific context sequentially until a final metastases prediction is made. b) Hidden unit (H_t) shared between steps (t) and receives the context of previous CT scan (x_t) and predicts the prognosis (y_t).

Autoencoder - The autoencoder [14] is a unsupervised algorithm, which means that unlike the previous supervised algorithms it does not require labelled data in the training process. The goal of autoencoders is to learn a compressed representation (code) of the input data by reconstructing it as the output of the network. By restricting the size of the code, the technique can discover the interesting structures of the data, and in the case of denoising autoencoder, even reconstruct noisy images. Depending on the characteristics of the input, the encoder and decoder can have different architectures, some based on multilayer perceptrons and other on convolutional neural networks.

Training process - The denoising autoencoder (Figure 6) contains an encoder which receives the noisy input, compresses into a small representation, called code, and is reconstructed by a decoder into the original noiseless input. Due to the small size of the code, the autoencoder learns the distinctive features of the image and learns to ignore random noise. During training, the weights of the neurons present in the encoder and decoder are continuously updated to reduce the difference between the original input and the output, called reconstruction error, to find useful patterns in the data.

One frequent use of denoising autoencoders is the extraction and compression of relevant features for the detection of genes correlated with the ER status of patients with breast cancer [15]. Figure 6 illustrates how the autoencoder is given a set of gene expression data with some noise with the task of compressing the data into an relevant representation (code).




Fig. 6: In a Denoising Autoencoder an encoder transforms a noisy gene expression data into a compressed representation (code) and the decoder transforms the code back into denoised version of original data.

III. INTERPRETABILITY CONCEPTS OVERVIEW: DESIDERATA, DIMENSIONS AND STRATEGIES

The significance of interpretability when developing ML solutions is well-known in academia and corporations. However, there is no consensus upon the definition of interpretability [16]. One of the most used definitions was presented by [17] which defined interpretability as the “ability to explain or to present in understandable terms to a human”, and will be used in this work.

A. Desiderata of Interpretability

The demand of interpretability arises due to a mismatch between the objectives of the model and of the users - clinicians and patients. Although DL techniques have reach human performance in melanoma diagnosis from dermoscopic images [4], or the detection of lymph node metastases in breast cancers from pathology images [5], the need to interpret them emerges, especially in healthcare contexts.

In addition to high accuracy of ML algorithms, users have additional desiderata. Doshi-Velez and Kim [17] specified five main desiderata for interpretability:

- **Fairness:** Assure that protected groups (e.g. gender, ethnicity) are not somehow discriminated against (explicit or implicit);
- **Privacy:** Assure that sensitive information is protected;
- **Reliability/Robustness:** Assure high algorithmic performance despite variation of parameter or input;
- **Causality:** Assure that the predicted change in output due to a perturbation will occur in the real system;
- **Trust:** Allow users to trust a system capable of explaining its decisions rather than a black box that just outputs the decision itself.

B. Dimensions of Interpretability

Interpretability methods can be characterized by a set of dimensions [18]: global and local interpretability, intrinsic and post-hoc interpretability and model-specific and model-agnostic interpretability. These will be described in what follows.

a) *Global and Local Interpretability:* This dimension reflects the scope of interpretability of a model and depicts the portion of predictions that the model can explain. To perform a classification task an ML algorithm first creates a data-driven model based on a set of input features (e.g.

age and sex) during the training phase. The objective of this phase is allowing neurons to select important features and learn relationships between them and the target output. Global interpretability aims to analyze this model, to understand the common patterns in the overall data that help make decisions, by studying the model's parameters (i.e. weights), and the learned relationships. Local interpretability aims to understand the relationship between the set of input features of a specific case and the model decision.

In our MLP example (Figure 2), based on the instances provided, the network learned relationships that help predict the tumor malignancy, based on its density, shape and breast side. As the breast side (left or right) where the tumor appears is not indicative of the level of malignancy, the network should have learned to discard this input feature.

Global interpretability could help understand which relationships the network learned, and for the example of breast side confirm that it was not used. Global interpretability can also help as know if non-random sources of noise which have been not been removed have affected the model's learning (e.g. artifacts). Local interpretability could help understand the importance of the input features in the malignancy prediction of a particular patient.

b) Intrinsic and Post-hoc Interpretability: While the increase of complexity of ANNs (i.e. number of neurons), help solve complex problems, it increase the difficulty to interpret them. Intrinsic interpretability refers to models which due to their simplicity are interpretable by themselves, such as decision trees or sparse linear models [18]. Complex models can increase their intrinsic interpretability by constraining their complexity or simplifying their behavior. Examples of these constraints are sparsity, monotonicity, adding domain knowledge, or even constraints on the complexity of the network by limiting the number of neurons or layers.

Post-hoc interpretability refers to the application of interpretability methods after the model's training [18]. Post-hoc methods help elucidate how the model works without constraining it.

In our MLP example, we could instead use a short decision tree or a small sparse MLP to achieve intrinsic interpretability or choose to maintain the complexity of the MLP and use a post-hoc method such as feature importance to understand the importance of the input features.

c) Model-specific and Model-agnostic: Another way to classify interpretability methods is based on the dependency the method has on the type of model which it tries to explain. Model-agnostic methods can be applied to different types of models, while model-specific methods are only applicable to a specific type of model [18].

In our example, while a model-agnostic method could extract the importance of the density and shape from a model trained from any ML algorithm, a model-specific method would only be able to do the same for similar models.

C. Interpretability Strategies

During the training phase, DL algorithms create data-driven models that can be interpreted using different strategies pro-

ducing different types of explanations. Namely feature importance, saliency map, model visualization, surrogate model, domain knowledge and example-based explanations, which will be introduced next.

1) Feature Importance: One of the more explored explanations is feature importance, which gives the importance or contribution of an input feature on the prediction of an example. Two main approaches are used for computing feature importance: sensitivity analysis [19] and decomposition [20], [21].

Sensitivity analysis computes the effects of the variation in the input variables in the model's output and help us answer the question "What change would make the instance more or less like a specific category?".

Decomposition approaches successively decomposes the importance of the output of a layer into previous layers, until the contribution that the input features have on the output is found. It help us answer the question "What was the feature's influence on the model's output?".

If we extract the feature importance of a decision of our example, it can have different meanings depending on the type of method used. High sensitivity values for density and shape means that their growth would also increase the prediction of malignancy. While high contribution values of density and shape means that the prediction of malignancy was highly influenced by the value of these features.

2) Saliency Map: When dealing with images, saliency maps [17] (or heatmaps) can be used to visually illustrate variations in the importance of different features, using color to convey the weight of pixel in a given prediction.

Similarly to feature importance, the pixel values of saliency maps can be obtained following two main approaches: Back-propagation methods compute the relevance of a pixel by propagating a signal from the output neuron backward through the layers to the input image in a single pass [21]. Perturbation methods compute pixel relevance by making small changes in the pixel value of the input image and compute how the changes affect the prediction [22].

An example of a saliency map, extracted from a CNN trained to predict the malignancy based on mammogram patches is seen in Figure 7. The red and yellow regions correspond to the most important regions of the image. The method correctly focus on the mass, supporting our confidence in the model's decisions.




Fig. 7: Example of a saliency map depicting the important pixels for malignancy prediction based on mammograms. Left: ground-truth expert segmentation. Right: saliency map, where the pixel intensity indicates the importance of the pixel in the classification.) [23].

3) *Model Visualization*: The ML algorithm receives an example with a set of input features, and in their internal process creates a combination of its features also called internal features. Some strategies help visualizing patterns detected in an image [24], whereas others help visualizing the feature distribution in the dataset [25], [26]. Also, whereas some strategies may ease to find the image that contains a pattern detected by the network [27], others artificially create images that accentuate the same patterns [28], [29].

In Figure 8 we can see regions of mammograms which contain patterns detected by individual filters of the CNN trained to diagnose the tumor malignancy.




Fig. 8: Illustration of the internal behaviour of a network unit by visualizing regions of mammograms with patterns detected by individual units of the network [27].

4) *Surrogate model*: A surrogate model are interpretable models that are trained to explain predictions of a black-box model. In the example of oncology, a rule list [30] can be extracted from a network allowing the clinician to understand the knowledge produced by the algorithm. Each rule specifies a condition which when evaluated as true produces on result (benign/malignant in malignancy diagnosis). One way of doing this is by creating a new dataset where each example of the dataset used to train the DL model is combined with its prediction and the task of the surrogate model is to predict this values.

While global surrogates models approximate the model in all the input space, local surrogate models approximate single predictions, which makes them more accurate and faithful to the model being explained.

To better understand what is a surrogate model, let's consider the example in Figure 9, where we can see a rule list extracted from a MLP that demonstrate its decisions. This surrogate model was built by iterating through the MLP neurons and inspecting the connections between the input features and the output label, so that they can be represented by rules. Decision tree is another appropriate type of surrogate model. These method could be seen as an unordered rule list where each leaf is a separate rule where the condition is the labels of the path from the root to the leaf.

5) *Domain Knowledge*: Although DL algorithms extract internal features (combination of input features) automatically during the training phase, the domain knowledge of the medical field which physicians have can be used to validate the decision of the network.

- Rule 1:** IF (density = 'high' or 'equal') and (shape = 'irregular') THEN **malignant**
- Rule 2:** IF (density = 'high') THEN **malignant**
- Rule 3:** IF (density = 'fat-containing') and (shape = 'irregular') THEN **benign**
- Rule 4:** OTHERWISE **benign**

Fig. 9: Rule list extracted from a MLP trained to predict the malignancy of a breast tumor using a surrogate model strategy.

The introduction of domain knowledge from medical doctors on training can help produce models that resemble how medical doctors diagnose or focus on the features or areas they pay particular attention to [31].

In the case of malignancy diagnosis, domain knowledge can be introduced directly as input feature, for example a discrete value indicating the shape of the tumor. Domain knowledge can also be used as an additional target variable (e.g. shape, density), besides malignancy, allowing to evaluate how well the model predicts both target variables similarly to how clinicians also take those variables into account.

6) *Example-based explanation*: Example-based explanation methods select examples of the dataset that explain the behavior of the network [18]. This behavior is usually explained using the internal features (combination of input features) extracted from the examples by the network.

Similar examples are examples of the dataset that have similar values on the internal features and produce the same prediction as the example whose prediction we are explaining [32].

Counterfactual explanations can be used to explain predictions of examples by finding small changes in the example that cause the network to change its prediction.

Usually examples of a dataset can be grouped together based on existing patterns. A prototype is a particular example of the dataset representative of its group.

Table I associates the interpretability strategies previously introduced with the dimensions of interpretability, namely scope and intrinsic vs. post-hoc. The dimensions of model specificity vs. agnostic was omitted as it depends on the actual algorithms used and not on the broader interpretability strategy.

Strategy	Dimensions	
	Scope	Intrinsic vs. Post-hoc
Feature Importance	Local	Post-hoc
Saliency Map	Local	Post-hoc
Model Visualization	Global	Post-hoc
Surrogate Model	Local/Global	Post-hoc
Domain Knowledge	Global	Intrinsic
Example-based	Global	Post-hoc

TABLE I: Association between interpretability strategies and dimensions of interpretability.

IV. INTERPRETING DEEP LEARNING IN ONCOLOGY

The use of DL techniques has become widespread in the oncology area, covering different pathologies, but their

interpretation remains an unexplored field [33], [34]. In this section, an overview of interpretability strategies applied to oncological diseases will be presented. The section will be divided into different diseases, namely breast cancer, skin cancer, lung cancer, brain cancer and other. This division was chosen to promote the best understanding of the area by the main target audience of this paper - oncologist, clinicians and other practitioners.

We conducted a search of papers in the PubMed database published between January 2014 and September 2020 with individual and combination of search terms such as “interpretability”, “deep learning”, “oncology”, “cancer” and “decision support systems”, and compiled the results in Table II. In total, 44 works were found, where the majority target in breast cancer (30%), skin cancer (23%), lung cancer (9%) and brain cancer (11%). The most common interpretability strategies were saliency maps (32%) and feature importance (20%) and among the prediction tasks, most works focused on diagnosis of malignancy (45%) and of different pathologies (27%).

Figure 10 helps visualize the distribution of papers based on different classifications present in Table II, namely the target disease and task as well as the interpretability strategy (explanation) and ANN technique (architecture).

A. Breast Cancer

Prediction of breast cancer malignancy has been one of the most successful applications of deep learning in oncology, achieving 87% sensitivity and 96% specificity when diagnosing mammograms [78]. It also is the main task on interpretability work (69% of breast cancer studies). Due to the availability of well-curated public datasets on breast cancer, mainly mammograms and hematoxylin and eosin (H&E) stained histological images, research in this area has taken a step forward.

When dealing with imaging data, researchers found it important to visualize the patterns detected by the networks either through model visualization techniques or with saliency maps, please refer to section III-C2. These patterns were then either validated by experts or correlated with medical concepts. For other types of data (e.g. gene expression, hand-crafted features), researchers mainly focused on computing feature importance or extracting surrogate models (i.e. rule lists). In what follows, we analyze with detail some of the main selected works on the topic.

Graziani *et al.* [35] visualized the patterns of a metastases detection CNN for WSI H&E images by synthesizing images that increase the network’s confidence on the prediction (Activation Maximization [28], [29]) and by extracting saliency maps [79]. They found that the network detected nuclei-resembling shapes and regions of nuclei with marked variations in size and irregular shapes. Hsieh *et al.* [27] used Network Dissection method [80] to visualize the patterns of individual filters of a malignancy classifier based on mammograms and developed a web-based tool which let experts label the patterns. Figure 11 shows an example of a pattern which was labeled as ‘Calcified Vessels’. Also, other BI-RADS [81]

medical concepts (e.g. mass margin) were found to overlap with patterns detected by the network.

Rather than being validated by experts, Graziani *et al.* [36], [37] introduced Regression Concept Vectors (an extension of Concept Activation Vectors [82]) which let them detect the importance of medical concepts (i.e. area, perimeter and contrast) on the decisions of a breast cancer malignancy classifier based on WSI H&E network, even though they were not present in the training dataset. Contrast was found to be positively correlated with malignancy, while correlation was negatively correlated. Kim *et al.* [38] used medical concepts during training, computing their importance alongside saliency maps to help explain the malignancy diagnosis of mammograms.

Antropova *et al.* [39] visualized the values of both deep features and hand-crafted features from different image modalities (i.e. Mammogram, Ultrasound, DCE-MRI) and found that their fusion improved malignancy diagnosis performance, most likely due to the low agreement between deep and handcrafted features.

Lee *et al.* [40] trained a malignancy diagnosis network able to justify its decisions both visually and textually. It trained a language model that composes text description [38], [76], [77], [83] from mammograms. Although the descriptions are still not sufficiently good (i.e. “There are sharp lines on some part of complexly formed mass.”), they show that this interpretability strategy has great potential.

When dealing with hand-crafted features relating with tumor size and shape, researchers found it important to simplify the network to behave linearly [41] making it easier to compute the feature importance, or extract simpler classifiers that could present physicians with simple rules (i.e. decision rules [42] and symbolic rules [43]) increasing interpretability.

Feature importance was the focus of most works dealing with gene expression data. For example, SALMON [44] predicted survival risk of patients with breast cancer, and feature importance of eigengene’s modules and other clinical information, they confirmed that age, progesterone receptor status and other five mRNA sequence data co-expression modules play pivotal roles in patient prognosis. Similar methods, using the H2O [84] library, were used to detect the important features in the detection of estrogen-receptor-positive (ER+) patients based on the classification of the Estrogen Receptor Status of breast cancer patients based on metabolomics data [45]. They found eight commonly enriched significant metabolomics pathways: isoleucine, putrescine, glycerol, 5'-deoxy-5'-methylthioadenosine, ornithine, tocopherol beta, phenylalanine, and arachidonic acid. Finally, Liu *et al.* [46] used an autoencoder to find clusters of breast cancer patients based on their gene expression and copy number alteration data, and visualized them using heatmaps. They found that the cluster of patients with ER-negative breast cancer patients usually have a poor prognosis.

B. Skin Cancer

Works in skin cancer almost evenly divided on the malignancy diagnosis and diagnosis of multiple skin diseases.

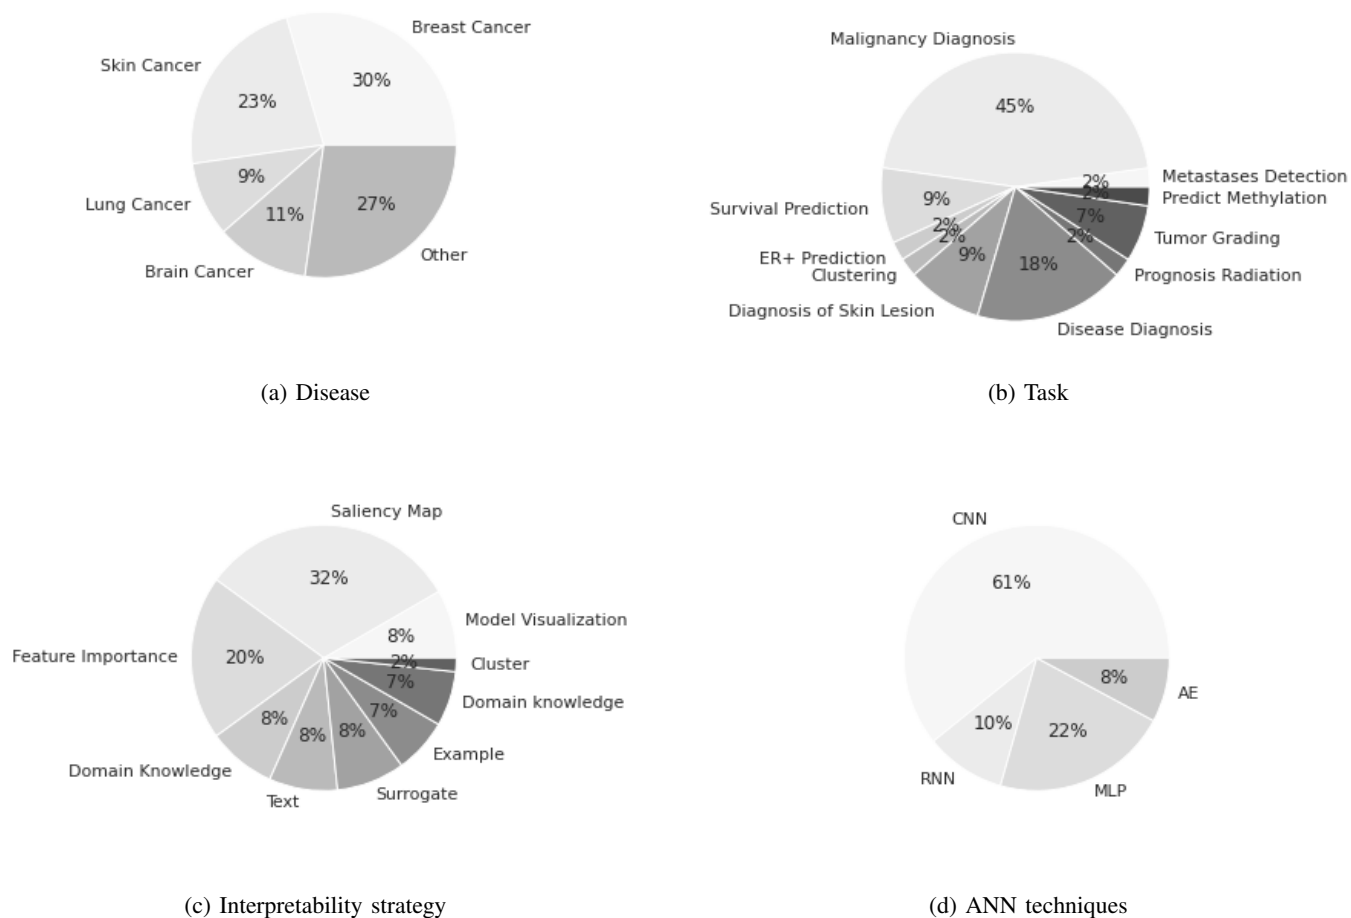


Fig. 10: Distribution of papers reviewed based on characteristics of Table II




Fig. 11: Example of pattern detected by the network and labeled by an expert as ‘Calcified Vessels’ in the web-based labeling tool [27].

The modality used was also divided between two types, dermoscopic images (70%) and H&E stained histopathological images (30%). Similarly to breast cancer detection, DL has also achieved great results in skin cancer detection based on medical imaging [85]. Interpretability methods for these pathologies ranged from saliency maps, model visualization,

rule extraction, text explanations and example-based explanations.

A simple visualization method was used to visualize the activation of neurons of a CNN trained to predict the malignancy of dermoscopic images [47]. Inspection of activations led to finding neurons related to medical concepts such as borders, lesions, and skin type, as well as different image artifacts such as hairs.

Cruz-Roa *et al.* [48] proposed a DL technique for the malignancy diagnosis using histological images and visualized the most salient patterns in that task which when validated by pathologists were found to be related large-dark nuclei. Researchers also tried to improve the quality of saliency maps by making changes on the architecture of the network when diagnosis malignancy based dermoscopic images [49] and diagnosis of skin diseases based on WSI H&E images [50]. PatchNet [49] found a trade-off between interpretability and performance, as smaller patch sizes provided saliency maps with better visual interpretability at the expense of worse generalization capabilities. Paschali *et al.* [50] also found that smaller convolutional filters resulted in more fine-grained saliency maps. Gonzalez-Diaz *et al.* [51] incorporated segmentation of lesion areas based on high-level dermoscopic

TABLE II: Summary of papers reviewed.

Ref	Disease	Task	Modality	Explanation	Architecture	Dataset
[35]	Breast Cancer	Metastases Detection	WSI H&E	Model Visualization, Saliency Map	CNN	Public
[27]	Breast Cancer	Malignancy Diagnosis	Mammogram	Model Visualization	CNN	Public
[36], [37]	Breast Cancer	Malignancy Diagnosis	WSI H&E	Feature Importance, Domain Knowledge	CNN	Public
[38]	Breast Cancer	Malignancy Diagnosis	Mammogram	Domain Knowledge, Saliency Map	CNN	Public
[39]	Breast Cancer	Malignancy Diagnosis	Mammogram, Ultrasound, MRI	Domain Knowledge	CNN	Public
[40]	Breast Cancer	Malignancy Diagnosis	Mammogram	Saliency Map, Text	CNN + RNN	Public
[41]	Breast Cancer	Malignancy Diagnosis	Hand-crafted	Feature Importance	CNN	Public
[42]	Breast Cancer	Malignancy Diagnosis	Hand-crafted from H&E	Surrogate	MLP	Private
[43]	Breast Cancer	Malignancy Diagnosis	Hand-crafted from H&E	Surrogate	MLP	Public
[44]	Breast Cancer	Survival Prediction	Gene expression, Biomarkers	Feature Importance	MLP	Public
[45]	Breast Cancer	ER+ Prediction	Metabolomics Data	Feature Importance	AE + MLP	Public
[46]	Breast Cancer	Clustering	Gene expression, CNA data	Model Visualization	AE	Public
[47]	Skin Cancer	Malignancy Diagnosis	Dermoscopic images	Model Visualization	CNN	Public
[48]	Skin Cancer	Malignancy Diagnosis	WSI H&E	Saliency Map	CNN	Private
[49]	Skin Cancer	Malignancy Diagnosis	Dermoscopic images	Saliency Map	CNN	Public
[50]	Skin Cancer	Diagnosis of Skin Lesion	WSI H&E	Saliency Map	CNN	Public
[51]	Skin Cancer	Diagnosis of Skin Lesion	Dermoscopic images	Saliency Map	CNN	Public
[52]	Skin Cancer	Malignancy Diagnosis	WSI H&E	Saliency Map	CNN	Public
[53]	Skin Cancer	Diagnosis of Skin Lesion	Dermoscopic images	Example	CNN	Public
[54], [55]	Skin Cancer	Malignancy Diagnosis	Dermoscopic images	Feature Importance, Example, Surrogate	MLP	Public
[56]	Skin Cancer	Diagnosis of Skin Lesion	Dermoscopic images	Example, Saliency Map	CNN	Public
[57]	Lung Cancer	Disease Diagnosis	Chest Radiograph	Saliency Map	CNN	Public
[58]	Lung Cancer	Malignancy Diagnosis	CT	Domain knowledge	CNN	Public
[59]	Lung Cancer	Malignancy Diagnosis	CT	Domain knowledge	CNN	Public
[60]	Lung Cancer	Prognosis Radiation	Biomarker, clinical data	Domain knowledge	AE + MLP	Private
[61]	Brain Cancer	Tumor Grading	MRI	Saliency Map	CNN	Public
[62]	Brain Cancer	Tumor Grading	MRI	Feature Importance, Saliency Map	MLP	Public
[63]	Brain Cancer	Predict Methylation State	MRI	Model Visualization	CNN + RNN	Public
[64]	Brain Cancer	Survival Prediction	MRI	Feature Importance	CNN	Public
[65]	Brain Cancer	Survival Prediction	WSI H&E, Biomarkers	Saliency Map	CNN	Public
[66]	Other	Malignancy Diagnosis	Gene expression	Feature Importance	MLP	Public
[67]	Other	Survival Prediction	Gene and protein expression	Feature Importance	MLP	Public
[68]	Other	Disease Diagnosis	RNA-seq expression, SVN data	Feature Importance, Surrogate	MLP	Private
[69]	Other	Disease Diagnosis	Volumetric Laser Endomicroscopy	Saliency Map	CNN	Private
[70]	Other	Disease Diagnosis	Endoscopic images	Saliency Map	CNN	Public
[71]	Other	Disease Diagnosis	WSI H&E	Saliency Map	CNN	Private
[72]	Other	Disease Diagnosis	DESI	Cluster	AE	Private
[73]	Other	Disease Diagnosis	Ophthalmic images	Domain Knowledge	CNN	Private
[74]	Other	Malignancy Diagnosis	Ultrasound	Domain knowledge	CNN	Private
[75]	Other	Malignancy Diagnosis	WSI H&E	Text, Saliency Map	CNN + RNN	Public
[76]	Other	Disease Diagnosis	Chest Radiograph	Text, Saliency Map, Text	CNN + RNN	Public
[77]	Other	Tumor Grading	WSI H&E	Text, Saliency Map	CNN + RNN	Private

features, and used these segmentations to diagnose of skin lesions and show relevant regions.

Example-based explanation are also useful interpretability strategies in skin cancer, as shown by Sadeghi *et al.* [53] which conducted a study which revealed that similar examples provided by DL techniques help users in classifying skin lesions from dermoscopic images. In the study, accuracy increased from 51% to 61% when the 15 most similar cases were provided to the users. Silva *et al.* [54], [55] unified complementary explanations to explain skin lesion predictions from dermoscopic images. The method extracted rules and presented them as text sentences alongside positive and a counter-factual examples for every decision. Also on the same task, Codella *et al.* [56] explained the decision with similar examples using k-nearest neighbors on the deep features and highlighted the most salient regions of the image.

C. Lung Cancer

Interpretability research on the diagnosis of lung cancer focused mainly on two modalities, Chest Radiography (X-Ray) or Computed Tomography (CT). Similarly, to breast and skin cancer, DT techniques have been shown to be able to reach human-level performance. In the diagnosis of 14 different pathologies from chest radiographs, a CNN achieved radiologist-level performance [57]. Radiologists confirmed, by inspecting saliency maps [86], that the network localizes accurately the lung masses.

Other works focused on the integration between hand-crafted features related to medical concepts and deep features. Paul *et al.* [58] developed a model for the malignancy diagnosis of lung cancer using CT images, and interpreted their correlation with medical features used by physicians by iteratively replacing deep features and evaluating the drop in confidence. Although deep features were not found to be perfectly correlated with medical features, they could represent 9 of the medical features with the deep features without losing performance. In the same task, Shen *et al.* [59] proposed to model that made high-level predictions for the tumor malignancy, and low-level predictions of medical features - calcification, subtlety, lobulation, sphericity, internal structure, margin, texture and spiculation. The approach achieved comparable or better results with state-of-the-art methods in the public Lung Image Database Consortium (LIDC).

Finally, Cui *et al.* [60] used a combination of hand-crafted features composed of clinical features and cancer biomarkers in a non-small cell lung cancer who received radiotherapy to predict the damage caused by the treatment. The results found that better performance was achieved by integrating the hand-crafted features with the deep features extracted from an autoencoder [87].

D. Brain Cancer

Unlike previous pathologies, brain cancer research deviates from diagnosis of diseases and focus on survival prediction

(40%) and tumor grading (40%), almost entirely based on Magnetic Resonance Imaging (MRI) (83%).

When performing tumor grading - distinguishing from lower grade gliomas from high grade gliomas from MRI - researchers have focused on producing saliency maps from the 3D MRI scans or Region of Interest (ROI) annotated by experts. Pereira *et al.* [61] extended existing saliency map methods for three dimensional inputs [79], [88]. The ROI classifier achieved better performance than the 3D scan (93% and 90% accuracy), but they were both able to locate the tumor. Pereira *et al.* [62] also used a feature importance method [89] to identify MRI sequences which were relevant for features extracted from the network, and then produce saliency maps. The sequences chosen were consistent with domain knowledge.

Han *et al.* [63] train a model to predict the methylation state of the MGMT regulatory regions using MRI of Glioblastoma Multiforme (GBM) patients, resulting in 62% accuracy. The MRI scans were extracted from the Cancer Imaging Archive (TCIA) [90] and the methylation data from the Cancer Genome Atlas (TCGA) [91]. The authors developed an online visualization tool which allows the user to load an MRI scan and visualize the activation of different filters. Through this the model was found to classify lesions with ring enhancement with negative methylation status and tumors with less clearly defined borders and heterogeneous texture with positive methylation status.

Lao *et al.* [64] constructed a model for survival prediction of patients with GBM based on deep features and hand-crafted features extracted from MRI. To reduce the number of features used, feature selection was done using feature importance methods to find features that were robust to tumor segmentation uncertainty, highly predictive and non-redundant. Survival prediction was also performed using histological samples and genomic data [65] with validation of produced saliency maps by expert pathologists.

E. Other Pathologies

Other oncological pathologies have been shown interested in interpretability using different modalities of data (not exclusively image). Researchers that applied DL techniques on data of multiple pathologies have sought to interpret them using feature importance. For example, Ahn *et al.* [66] trained a network for malignancy diagnosis based on gene-expression data from multiple tissues and by computing the feature importance of individual genes on the diagnosis found a subgroup suspected to be oncogene-addicted as an individual gene contributes extensively in the classification. Similarly, Yousefi *et al.* [67] proposed a model for the survival prediction based on clinical, gene-expression and protein-expression data of multiple tissues and computed the sensitivity of each feature on the survival risk, identifying that TGF-Beta 1 signaling and epithelial-mesenchymal transition (EMT) gene sets are associated with poor prognosis. Oni *et al.* [68] diagnosed eight different cancer types from RNA-seq expression and single nucleotide variation (SNV) data. To explain its decisions, a linear surrogate model [89] was extracted, where its coefficient's magnitude corresponded to importance of the genes

in the prediction. The location and variability of explanations were visualized using 2D embeddings of the RNA-seq input data. They found genes related to cell proliferation and tumor growth were important for the diagnosis.

In the diagnosis of early Barrett's Neoplasia using Volumetric Laser Endomicroscopy [69], saliency maps [86] focused on the glands located around the first layers of the esophagus in high-grade dysplasia cases, and on homogeneous esophagus layers in non-dysplastic Barrett's esophagus cases. Garcia-Peraza-Herrera *et al.* [70] extended the same saliency map method to interpret the diagnosis of esophageal cancer based on endoscopic images. By computing saliency maps of different resolutions they were able to detect unhealthy patterns and diseased tissue.

Korbar *et al.* [71] interpreted the diagnosis of colorectal polyps based on histological images using saliency maps [79], [86] and found that by adding a boundary box around them increased their similarity with pathologists' segmentations.

Inglese *et al.* [72] used DL techniques to find a high-level representation of mass spectrometry imaging data from colorectal adenocarcinoma biopsies. The features extracted from the network were visualized in two dimensions using t-SNE [92] unveiling clusters with different chemical and biological interactions occurring.

Zhang *et al.* [73] developed a diagnostic system of ophthalmic images that explained the diagnosis with sub-tasks. In addition to the diagnosis disease, the network segmented important anatomical regions, and detected other illnesses. The results show an accuracy of 93% on the diagnosis, localization accuracy of the foci of 82% in normal lighted images and 90% in fluorescein sodium eye drops.

Zhang *et al.* [74] proposed a system for diagnosing the malignancy of thyroid nodules on ultrasound with performance comparable with radiologists. The network provides prediction on medical concepts based on the TI-RADS lexicon.

The automatic generation of text reports based on medical imaging system is also an active research area. Zhang *et al.* [75] presented network trained on H&E patches for the malignancy diagnosis of bladder cancer, and conditioned a RNN-based language model to generate text descriptions and visual attention (i.e. saliency maps) highlighting regions of the image relevant for specific parts of the text (Figure 12). Similarly, TieNet [76] provided the same explanation for the network which diagnoses diseases based on chest radiographs and generates text descriptions with similar visual attention. MDNet [77] establishes a relationship between histological images of bladder cancer and diagnostic reports to generate text descriptions and provide visual attention for specific parts of the text.

V. OPEN ISSUES AND PROMISING RESEARCH DIRECTIONS

As DL grows in popularity, so does the need for interpretability in the dichotomy between ML and medical practice. From this survey, it becomes clear that there are four main issues that need more attention: (1) limitation on the applications of interpretability methods; (2) limitation on medical tasks explored; (3) lack of reliability of some interpretability methods;

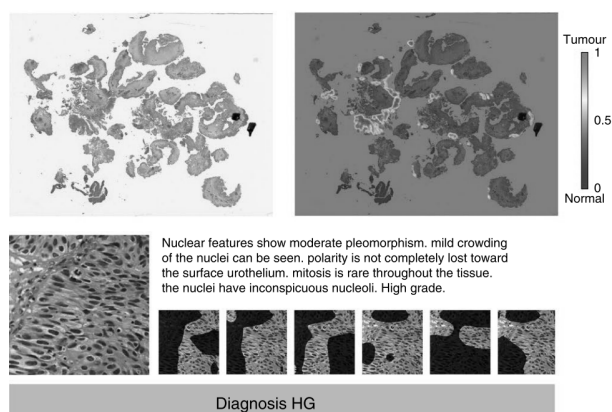


Fig. 12: Left: H&E stained whole-slide tissue image. Right: saliency map generated [75]. Bottom: description generated for the image and feature-aware attention maps.

and (4) lack of evaluation metrics for interpretability methods. Throughout this section, we will provide a discussion on the former four issues.

A. Limitation on the applications of interpretability methods

Du *et al.* [93] classified three major application of interpretability strategies: model validation, model debugging and knowledge discovery.

Model validation verifies that the model was able to learn useful knowledge and avoid learning bias information. The majority of works reviewed follow in this category, for example works which explored the use of saliency maps mainly focused on verifying that the region highlighted corresponded to regions segmented by experts.

Other applications for the interpretation of deep learning models, such as model debugging and knowledge discovery, were overlooked by the current literature and constitute promising directions to further improve the diagnostic capabilities of models and discover new insights on the biology of different cancer diseases.

Model debugging aims at analyzing what leads to the misbehaviour of models and erroneous predictions. Interpretability can help to uncover the reason for this misbehaviour, by inspecting the examples what were misclassified by the model, examples that have artifacts from the data collection (e.g. metal tools in a CT scan, hairs in a dermoscopic image), in addition to difficult to diagnose cases. Model debugging is also extremely relevant when generalizing the model for other hospital data or for clinical use where the risk for misbehavior is much bigger. This application is still overlooked in current works in the field of oncology.

Carlini *et al.* [94] demonstrated that standard models can make perfect predictions in random training set while performing poorly on the test set. This proves the model's ability to memorize the input data even if it is random which causes low generalization to unseen data. The lack of generalization of models which can be caused by overfitting to the training dataset must be an active concern of all ML practitioners, especially deep learning techniques as the high complexity of

the models coupled with an low data size increases risk of overfitting.

Another issue related with model debugging is adversarial attacks which consist on inputs that are intentionally crafted to force the model to make a mistake. Finlayson [95] demonstrated how an adversarial noise added to a dermoscopic image previously diagnosed as benign with over 99% confidence by a highly accurate model resulted with the model predicting malignant with 100% of confidence even though the difference is imperceptible to the human eye. Finlayson [95] also pointed at insurance claims approvals as a possible motivation for adversarial attacks.

Another problem with generalizability is discriminatory bias where models learn unintended associations regarding minority subgroups due to bias in the data used to train the model [96]. An example is how malignancy diagnosis systems with accuracy similar to that of board-certified dermatologists under-performed on images of lesions in skin of color due to the majority of training examples represent fair skinned patients [96].

Discriminatory bias is not the only type of bias which can cause problems as there have been several instances where exceptional results have been obtained from the model learning to distinguish slides based on the hospital they came from or the clinicians that generated the ground truth rather than actual evidence in the slide [96]. For example, a system for the detection of pneumonia on chest x-rays was able learn to associate the use of a portable x-ray machine with pneumonia [97].

Knowledge discovery allows physicians and researchers to obtain new insights on the physiology of the disease by interpreting the deep learning model and its decision process, such as finding that HER2 receptor over-expression is related to breast cancer. Knowledge discovery could lead to finding other receptors help in the characterization of cancer diseases that are still unknown to this date. While some visualization methods have been used to discover cluster of patients with specific characteristics [46], [72], this direction of research is still mostly unexplored.

B. Limitation on medical tasks explored

Analysis on the results of the review (Figure 10) shows that 72% of works focus on some type of disease classification (45% malignancy diagnosis, 18% disease diagnosis, 9% diagnosis of skin lesion). This shows a great imbalance as their exists many more medical tasks in the oncology field with promising results but still lack interpretability. In the following sub-section relevant work on other medical tasks will be briefly reviewed. Those medical tasks are:

- Tumor or lesion segmentation: identify the set of voxels which make up the lesions or tumors present [98], [99];
- Organ and substructure segmentation: identify the set of voxels which make up either the contour or the interior of the objects of interest [100];
- Cancer prognosis: estimate the likely course and outcome of a disease [101], [102];

- Radiation treatment planning: determine location and dosage to deliver the most desirable dose distribution of radiotherapy [103];
- Image registration: seeks to determine a transformation that will map two volumes (source and reference) to the same coordinate system [104];
- Image generation and enhancement: includes many different tasks to improve quality of the input from removing obstructing artifacts or noise in images to complete missing data [105]–[107].

Tumor and lesion segmentation is an important first step for numerous other tasks such as diagnosis and treatment planning, in order to evaluate the extend of the diseased tissue. DL techniques have achieved state-of-the-art results in brain tumor segmentation from MRI scans [98], [99]. The same type of networks have also been used in the segmentation of different lesion of the skin based on dermoscopic images [108], [109].

Segmentation of organs and substructures is also an critical step before radiotherapy in order to decide the which regions to avoid targeting with radiation. One example of it is the segmentation of organs from abdominal CT scans [100].

Cancer prognosis is comprised of a large number of sub-tasks such as survival prediction and prediction of likelihood of metastases. Zhu *et al.* [101] for example, reviewed a large number of studies which applied DL techniques to different cancer prognosis tasks such as cancer recurrence, progression and survival prediction [101]. Other studies focused on sub-tasks which concern with the progression of the disease after treatment, from the prediction of future distant metastases and local-regional recurrence using pre-treatment, post-treatment and follow-up medical imaging scans [102].

Radiation treatment planning requires not only the segmentation of diseased tissue but also the dosage that should be used. An CNN-based model was used to MRI to accurately transfer contrast into CT images with clearly identified air, brain soft tissue, and bone highly similar to that of current methods based on CT and used in medical practice [103].

Image registration, also known as image fusion, is commonly used to combine two modalities - for example PET-CT is obtained by combining two different modalities (PET and CT), but also multiple images of the same modalities. Fu *et al.* [104] review a large number of DL techniques proposed for the image registration of different modalities such as T1 and T2 MRIs and MRI and CT.

In addition, DL approaches also has seen success in restoring medical images corrupted with noise or artifacts, but the interpretation of the reasoning behind this process has also been pointed out as a challenge [110]. The extensive use of CT in medical analysis has raise some concerns due to the large dose of radiation that it delivers to the patient. Low dose CT is a solution for this problem, but by using lower radiation amounts, noise and artifacts become a problem. DL techniques have been proposed to reconstruct low dose CT images and recover from noise and streaking artifacts caused by metal objects [105]–[107].

Even though DL techniques have help the numerous problems pointed out above, they all face the same obstacle which

prevents their use in clinical practice, the lack of interpretability. Future research efforts should then be targeted in the exploration of other application of interpretability methods other than model validation and different cancer tasks than disease diagnosis. With the expansion of cancer applications, other interpretability strategies will emerge based on images (most used modality) and other modalities that may be more associated with other problems.

C. Lack of reliability of some interpretability methods

Some post-hoc interpretability methods can present bias [111], [112] and might not be representative of the behavior of the model they are trying to explain [113]. This happens because although explanations should approximate as much as possible the actual behaviour of the model, during the process of optimization (e.g. backpropagation) some inputs given to the network are outside the distribution of the training data and can trigger artifacts of the deep learning model.

As different interpretation methods sometimes focus on distinct aspects of the model [107], a promising direction to improve the reliability of the interpretations is deploy an ensemble of complementary interpretability methods. Furthermore, interpretability methods should also be provided with imperfect data (i.e. noisy) to guarantee robustness to noise.

D. Lack of evaluation metrics for interpretability methods

To quantitatively evaluate an interpretability method without the validation of an expert requires a formal definition of interpretability and the use of a proxy metric describing the quality of the explanation [17]. The lack of ground-truth explanations, for example the expert annotated tumor segmentations which indicated what the expected value of a saliency map should be, makes it difficult to make quantitative analysis of the results, and generalize the obtained results. One of possible solutions to solve this issue is to conduct a comparison study between the interpretation produced by the deep model and one produced by a set of physicians. However, and once again, this solution may not be generalizable, hence most studies conducting evaluation by letting experts (e.g. pathologist) compare the explanations of few number of selected examples and their domain knowledge.

Future research should help find interpretability metrics able to assess methods based on three factors. First, evaluate how faithful the explanations are to the actual model's behaviour. Second, evaluate how easily the explanations are understood by the physician. Third, evaluate the usefulness of the explanation of its target application (i.e. model validation). Only by evaluating these factors can explanations extracted from deep learning models be truly trusted and applied in clinical practice.

VI. CONCLUSIONS

Interpretability of deep learning is a growing field with mostly open problems and many opportunities for the field of medicine and oncology.

The lack of interpretability in deep learning has been pointed out as a major problem by many researchers that have

studied the application of deep learning in various areas of medicine and bioinformatics [33], [34], [114].

In this work, we presented an easy guide for oncologists where we introduced various deep learning techniques and illustrated how the decisions of these could be interpreted with self-explanatory oncological cases to better illustrate. We also review the related research on the application of interpretability methods for cancer diseases, summarizing their main conclusions.

To the extent of the authors' knowledge, such comprehensive review on the interpretability of DL models for cancer diseases has not been previously performed. Overall, a high number of studies focused on breast, skin and brain cancers (60%) and on the explanation of the importance of tumor characteristics like tumor dimensions and shape, in the prediction of decision system. The majority of DL techniques interpreted were multilayer perceptrons and convolutional neural networks, often used to predict based on raw images or handcrafted features extracted from them.

As discussed in the previous section, three main issues were identified: (1) limitation on the applications of interpretability methods; (2) lack of reliability of some interpretability method; and (3) lack of evaluation metrics for interpretability methods.

Future research should go beyond model validation and apply interpretability to understand how models misbehave, as well as discover new knowledge about different cancer diseases. Also, although DL has been successful in many cancer tasks (e.g. tumor segmentation, cancer prognosis and image registration), works aim at interpreting models on these tasks remain unexplored. Lastly, future research in the design of evaluation metrics and frameworks is mandatory to assess the reliability of AI systems and for increasing the trust to be used on clinical practice.

CONFLICT OF INTEREST

The authors declare no competing interests.

REFERENCES

[1] Futurist, "FDA Approvals For Smart Algorithms In Medicine In One Giant Infographic," <https://medicalfuturist.com/fda-approvals-for-algorithms-in-medicine/>, 2019, accessed: 2019-12-01.

[2] W. Penny and D. Frost, "Neural networks in clinical medicine," *Medical decision making : an international journal of the Society for Medical Decision Making*, vol. 16, pp. 386–98, 10 1996.

[3] F. Rosenblatt, *The Perceptron, a Perceiving and Recognizing Automaton Project Para*, ser. Report: Cornell Aeronautical Laboratory. Cornell Aeronautical Laboratory, 1957.

[4] T. J. Brinker, A. Hekler, A. H. Enk, J. Klode, A. Hauschild, C. Berking, Schilling, and et al, "A convolutional neural network trained with dermoscopic images performed on par with 145 dermatologists in a clinical melanoma image classification task," *European Journal of Cancer*, vol. 111, pp. 148–154, 4 2019.

[5] B. E. Bejnordi and et al., "Diagnostic assessment of deep learning algorithms for detection of lymph node metastases in women with breast cancer," *Journal of the American Medical Association (JAMA)*, vol. 318, no. 22, pp. 2199–2210, 2017.

[6] M. Reyes, R. Meier, S. Pereira, C. Silva, M. Dahlweid, H. Tengg-Kobligk, R. Summers, and R. Wiest, "On the interpretability of artificial intelligence in radiology: Challenges and opportunities," *Radiology: Artificial Intelligence*, vol. 2, p. e190043, 05 2020.

[7] S. Liu, X. Wang, M. Liu, and J. Zhu, "Towards better analysis of machine learning models: A visual analytics perspective," *Visual Informatics*, vol. 1, no. 1, pp. 48–56, 2017.

[8] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, "Learning representations by back-propagating errors," *Nature*, vol. 323, no. 6088, pp. 533–536, 10 1986.

[9] F. Shahidi, S. Mohd Daud, H. Abas, N. A. Ahmad, and N. Maarop, "Breast cancer classification using deep learning approaches and histopathology image: A comparison study," *IEEE Access*, vol. 8, pp. 187 531–187 552, 2020.

[10] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278–2323, 1998.

[11] Y. Bengio, "Deep Learning of Representations for Unsupervised and Transfer Learning," in *JMLR: Workshop and Conference Proceedings*, vol. 7, 2011, pp. 1–20.

[12] M. H. Hesamian, W. Jia, X. He, and P. Kennedy, "Deep learning techniques for medical image segmentation: Achievements and challenges," *Journal of digital imaging*, vol. 32, no. 4, pp. 582–596, Aug 2019, 31144149[pmid]. [Online]. Available: <https://pubmed.ncbi.nlm.nih.gov/31144149>

[13] Y. Xu, A. Hosny, R. Zeleznik, C. Parmar, T. Coroller, I. Franco, R. H. Mak, and H. J. Aerts, "Deep learning predicts lung cancer treatment response from serial medical imaging," *Clinical Cancer Research*, vol. 25, no. 11, pp. 3266–3275, 2019. [Online]. Available: <https://clincancerres.aacrjournals.org/content/25/11/3266>

[14] P. Baldi, "Autoencoders, unsupervised learning and deep architectures," in *Proceedings of the 2011 International Conference on Unsupervised and Transfer Learning Workshop - Volume 27*, ser. UTLW'11. JMLR.org, 2011, p. 37–50.

[15] J. Tan, M. Ung, C. Cheng, and C. Greene, "Unsupervised feature construction and knowledge extraction from genome-wide assays of breast cancer with denoising autoencoders," *Pacific Symposium on Biocomputing. Pacific Symposium on Biocomputing*, vol. 20, pp. 132–43, 01 2015.

[16] R. Guidotti, A. Monreale, S. Ruggieri, F. Turini, F. Giannotti, and D. Pedreschi, "A Survey of Methods for Explaining Black Box Models," *ACM Computing Surveys*, vol. 51, no. 5, pp. 1–42, 08 2018.

[17] F. Doshi-Velez and B. Kim, "Towards A Rigorous Science of Interpretable Machine Learning," *ArXiv e-prints*, feb 2017.

[18] C. Molnar, "Interpretable Machine Learning," 2019, <https://christophm.github.io/interpretable-ml-book/>.

[19] D. Baehrens, T. Schroeter, S. Harmeling, M. Kawanabe, K. Hansen, and K.-R. Müller, "How to Explain Individual Classification Decisions," *Journal of Machine Learning Research*, vol. 11, pp. 1803–1831, 08 2010.

[20] S. Bazen and X. Joutard, "The taylor decomposition: A unified generalization of the oxaca method to nonlinear models," 05 2013.

[21] S. Bach, A. Binder, G. Montavon, F. Klauschen, K.-R. Müller, and W. Samek, "On Pixel-Wise Explanations for Non-Linear Classifier Decisions by Layer-Wise Relevance Propagation," *PLOS ONE*, vol. 10, no. 7, pp. 1–46, 07 2015.

[22] K. Simonyan, A. Vedaldi, and A. Zisserman, "Deep Inside Convolutional Networks: Visualising Image Classification Models and Saliency Maps," *ArXiv e-prints*, 2013.

[23] K. J. Geras, R. M. Mann, and L. Moy, "Artificial intelligence for mammography and digital breast tomosynthesis: Current concepts and future perspectives," *Radiology*, vol. 293, no. 2, p. 246–259, November 2019. [Online]. Available: <https://europepmc.org/articles/PMC6822772>

[24] J. Yosinski, J. Clune, A. Nguyen, T. Fuchs, and H. Lipson, "Understanding Neural Networks Through Deep Visualization," in *International Conference on Machine Learning - Deep Learning Workshop*, jun 2015, p. 12.

[25] L. McInnes, J. Healy, and J. Melville, "Umap: Uniform manifold approximation and projection for dimension reduction," *ArXiv e-prints*, 2018.

[26] M. Graziani, J. M. Brown, V. Andrearczyk, V. Yildiz, J. P. Campbell, D. Erdogmus, S. Ioannidis, M. F. Chiang, J. Kalpathy-Cramer, and H. Muller, "Improved interpretability for computer-aided severity assessment of retinopathy of prematurity," in *Proceedings of Computer-Aided Diagnosis*, 2019, p. 63.

[27] S. Hsieh, C. D. Lehman, V. Dialani, B. Zhou, D. Peck, G. Patterson, L. Mackey, V. Syrgkanis, and J. Wu, "Expert identification of visual primitives used by CNNs during mammogram classification," in *Proceedings of Computer-Aided Diagnosis*, 2018, p. 100.

[28] A. Nguyen, A. Dosovitskiy, J. Yosinski, T. Brox, and J. Clune, "Synthesizing the Preferred Inputs for Neurons in Neural Networks via Deep Generator Networks," in *Proceedings of the 30th International*

- Conference on Neural Information Processing Systems*, ser. NIPS'16. Curran Associates Inc., 2016, pp. 3395–3403.
- [29] C. Olah, A. Mordvintsev, and L. Schubert, "Feature Visualization," *Distill*, 2017, <https://distill.pub/2017/feature-visualization>.
- [30] M. T. Ribeiro, S. Singh, and C. Guestrin, "Anchors: High-precision model-agnostic explanations," in *AAAI*, 2018.
- [31] X. Xie, J. Niu, X. Liu, Z. Chen, S. Tang, and S. Yu, "A survey on incorporating domain knowledge into deep learning for medical image analysis," *Medical Image Analysis*, vol. 69, p. 101985, Apr 2021. [Online]. Available: <http://dx.doi.org/10.1016/j.media.2021.101985>
- [32] R. Caruana, H. Kangaroo, J. Dionisio, U. Sinha, and D. Johnson, "Case-based explanation of non-case-based learning methods," in *Proceedings of the AMIA Symposium*, 02 1999, pp. 212–5.
- [33] A. Vellido, "The importance of interpretability and visualization in machine learning for applications in medicine and health care," *Neural Computing and Applications*, vol. 0123456789, pp. 1–15, feb 2019. [Online]. Available: <https://doi.org/10.1007/s00521-019-04051-w>
- [34] F. Cabitza, R. Rasoini, and G. F. Gensini, "Unintended Consequences of Machine Learning in Medicine," *JAMA*, vol. 318, no. 6, pp. 517–518, 08 2017. [Online]. Available: <https://doi.org/10.1001/jama.2017.7797>
- [35] M. Graziani, V. Andrearczyk, and H. Müller, "Visual interpretability for patch-based classification of breast cancer histopathology images," in *Proceedings of Medical Imaging with Deep Learning*, 2018, pp. 1–4.
- [36] M. Graziani, V. Andrearczyk, and H. Müller, "Regression Concept Vectors for Bidirectional Explanations in Histopathology," in *Understanding and Interpreting Machine Learning in Medical Image Computing Applications*. Springer International Publishing, 2018, pp. 124–132.
- [37] M. Graziani, V. Andrearczyk, S. Marchand-maillet, and H. Müller, "Concept attribution: Explaining CNN decisions to physicians," *Computers in Biology and Medicine*, vol. 123, p. 103865, 2020. [Online]. Available: <https://doi.org/10.1016/j.combiomed.2020.103865>
- [38] S. T. Kim, J. H. Lee, H. Lee, and Y. M. Ro, "Visually interpretable deep network for diagnosis of breast masses on mammograms," *Physics in Medicine and Biology*, vol. 63, no. 23, dec 2018.
- [39] N. Antropova, B. Huynh, and M. Giger, "A deep feature fusion methodology for breast cancer diagnosis demonstrated on three imaging modality datasets," *Medical Physics*, vol. 44, 07 2017.
- [40] H. Lee, S. T. Kim, and Y. M. Ro, "Generation of Multimodal Justification Using Visual Word Constraint Model for Explainable Computer-Aided Diagnosis," *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 11797 LNCS, pp. 21–29, oct 2019. [Online]. Available: <http://arxiv.org/abs/1906.03922>
- [41] D. Alvarez-Melis and T. Jaakkola, "Towards Robust Interpretability with Self-Explaining Neural Networks," in *Advances in Neural Information Processing Systems 31*, S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, Eds. Curran Associates, Inc., 2018, pp. 7775–7784.
- [42] J. P. Amorim, I. Domingues, P. Abreu, and J. Santos, "Interpreting deep learning models for ordinal problems," in *ESANN*, 2018, pp. 373–377.
- [43] G. Bologna and Y. Hayashi, "Characterization of symbolic rules embedded in deep DIMLP networks: A challenge to transparency of deep learning," *Journal of Artificial Intelligence and Soft Computing Research*, vol. 7, no. 4, pp. 265–286, 2017.
- [44] Z. Huang, X. Zhan, S. Xiang, T. S. Johnson, B. Helm, C. Y. Yu, J. Zhang, P. Salama, M. Rizkalla, Z. Han, and K. Huang, "Salmon: Survival analysis learning with multi-omics neural networks on breast cancer," *Frontiers in Genetics*, vol. 10, no. 4, pp. 1–13, 2019.
- [45] F. M. Alakwaa, K. Chaudhary, L. X. Garmire, and B. G. Program, "Deep Learning Accurately Predicts Estrogen Receptor Status in Breast Cancer Metabolomics Data," *Journal of Proteome Research*, pp. 337–347, 2018.
- [46] Q. Liu and P. Hu, "Association analysis of deep genomic features extracted by denoising autoencoders in breast cancer," *Cancers*, vol. 11, no. 4, 2019.
- [47] P. Van Molle, M. De Strooper, T. Verbelen, B. Vankeirsbilck, P. Simoens, and B. Dhoedt, "Visualizing Convolutional Neural Networks to Improve Decision Support for Skin Lesion Classification," in *Understanding and Interpreting Machine Learning in Medical Image Computing Applications*. Springer International Publishing, 2018, pp. 115–123.
- [48] A. Cruz-Roa and et al., "Automatic detection of invasive ductal carcinoma in whole slide images with convolutional neural networks," in *Medical Imaging: Digital Pathology*, vol. 9041. International Society for Optics and Photonics, 4 2014, p. 904103.
- [49] A. Radhakrishnan, C. Durham, A. Soylemezoglu, and C. Uhler, "Patchnet: Interpretable Neural Networks for Image Classification," *ArXiv e-prints*, 5 2017.
- [50] M. Paschali, M. Ferjad Naeem, W. Simson, K. Steiger, M. Mollenhauer, N. Navab, M. F. Naeem, W. Simson, K. Steiger, M. Mollenhauer, and N. Navab, "Deep Learning Under the Microscope: Improving the Interpretability of Medical Imaging Neural Networks," *ArXiv e-prints*, pp. 1–9, 2019. [Online]. Available: <http://arxiv.org/abs/1904.03127>
- [51] I. Gonzalez Diaz, "DermaKNet: Incorporating the knowledge of dermatologists to Convolutional Neural Networks for skin lesion diagnosis," *IEEE Journal of Biomedical and Health Informatics*, vol. 2194, pp. 1–14, 2018.
- [52] A. A. Cruz-Roa, J. E. Arevalo Ovalle, A. Madabhushi, and F. A. González Osorio, "A Deep Learning Architecture for Image Representation, Visual Interpretability and Automated Basal-Cell Carcinoma Cancer Detection," in *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. Springer Berlin Heidelberg, 2013, pp. 403–410.
- [53] M. Sadeghi, P. K. Chilana, and M. S. Atkins, "How Users Perceive Content-Based Image Retrieval for Identifying Skin Images," in *Understanding and Interpreting Machine Learning in Medical Image Computing Applications*. Springer International Publishing, 2018, pp. 141–148.
- [54] W. Silva, K. Fernandes, M. J. Cardoso, and J. S. Cardoso, "Towards Complementary Explanations Using Deep Neural Networks," in *Understanding and Interpreting Machine Learning in Medical Image Computing Applications*. Springer International Publishing, 2018, pp. 133–140.
- [55] W. Silva, K. Fernandes, and J. S. Cardoso, "How to produce complementary explanations using an ensemble model," in *2019 International Joint Conference on Neural Networks (IJCNN)*, 2019, pp. 1–8.
- [56] N. C. F. Codella, C.-C. Chung-Ching Lin, A. Halpern, M. Hind, R. Feris, and J. R. Smith, "Collaborative Human-AI (CHAI): Evidence-Based Interpretable Melanoma Classification in Dermoscopic Images," in *Understanding and Interpreting Machine Learning in Medical Image Computing Applications*. Springer International Publishing, 2018, pp. 97–105.
- [57] P. Rajpurkar and et al., "Deep learning for chest radiograph diagnosis: A retrospective comparison of the CheXNeXt algorithm to practicing radiologists," *PLOS Medicine*, vol. 15, no. 11, pp. 1–17, 11 2018.
- [58] R. Paul, Y. Liu, Q. Li, L. Hall, and D. Goldgof, "Representation of Deep Features using Radiologist defined Semantic Features," in *Proceedings of the International Joint Conference on Neural Networks (IJCNN)*, 11 2018, pp. 1429–1435.
- [59] S. Shen, S. X. Han, D. R. Aberle, A. A. T. Bui, and W. Hsu, "An Interpretable Deep Hierarchical Semantic Convolutional Neural Network for Lung Nodule Malignancy Classification," *ArXiv e-prints*, 2018.
- [60] S. Cui, Y. Luo, H.-H. Tseng, R. Ten Haken, and I. El Naqa, "Combining handcrafted features with latent variables in machine learning for prediction of radiation-induced lung damage," *Medical Physics*, vol. 46, 03 2019.
- [61] S. Pereira, R. Meier, V. Alves, M. Reyes, and 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, 2018, pp. 106–114.
- [62] S. Pereira, R. Meier, R. McKinley, R. Wiest, V. Alves, C. A. Silva, and M. Reyes, "Enhancing interpretability of automatically extracted machine learning features: application to a RBM-Random Forest system on brain lesion segmentation," *Medical Image Analysis*, vol. 44, pp. 228–244, 02 2018.
- [63] L. Han and M. R. Kamdar, "MRI to MGMT: predicting methylation status in glioblastoma patients using convolutional recurrent neural networks," *Analytical Chemistry*, vol. 25, no. 4, pp. 368–379, 2015.
- [64] J. Lao, Y. Chen, Z. C. Li, Q. Li, J. Zhang, J. Liu, and G. Zhai, "A Deep Learning-Based Radiomics Model for Prediction of Survival in Glioblastoma Multiforme," *Scientific Reports*, vol. 7, no. 1, dec 2017.
- [65] P. Mobadersany, S. Yousefi, M. Amgad, D. A. Gutman, J. S. Barnholtz-Sloan, J. E. Velázquez Vega, D. J. Brat, and L. A. D. Cooper, "Predicting cancer outcomes from histology and genomics using convolutional networks," *Proceedings of the National Academy of Sciences*, vol. 115, no. 13, pp. E2970–E2979, 2018.
- [66] T. Ahn, T. Goo, C. H. Lee, S. Kim, K. Han, S. Park, and T. Park, "Deep Learning-based Identification of Cancer or Normal Tissue using Gene Expression Data," *Proceedings of the IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, pp. 1748–1752, 2019.

- [67] S. Yousefi, F. Amrollahi, M. Amgad, C. Dong, J. E. Lewis, C. Song, D. A. Gutman, S. H. Halani, J. Enrique Velazquez Vega, D. J. Brat, and L. A. D. Cooper, "Predicting clinical outcomes from large scale cancer genomic profiles with deep survival models," *Scientific Reports*, vol. 7, p. 11707, 2017.
- [68] O. Oni and S. Qiao, "Model-Agnostic Interpretation of Cancer Classification with Multi-Platform Genomic Data," in *Proceedings of the 10th ACM International Conference on Bioinformatics, Computational Biology and Health Informatics*, ser. BCB '19. New York, NY, USA: ACM, 2019, pp. 34–41.
- [69] R. Fonollà and et al., "Ensemble of Deep Convolutional Neural Networks for Classification of Early Barrett's Neoplasia Using Volumetric Laser Endomicroscopy," *Applied Sciences*, vol. 9, no. 11, p. 2183, 2019.
- [70] L. C. Garcia-Peraza-Herrera, M. Everson, W. Li, I. Luengo, L. Berger, O. Ahmad, L. Lovat, H.-P. Wang, W.-L. Wang, R. Haidry, D. Stoyanov, T. Vercauteren, and S. Ourselin, "Interpretable Fully Convolutional Classification of Intrapapillary Capillary Loops for Real-Time Detection of Early Squamous Neoplasia," *ArXiv e-prints*, pp. 1–8, 2018.
- [71] B. Korbar, A. M. Olofson, A. P. Mirafior, C. M. Nicka, M. A. Suriawinata, L. Torresani, A. A. Suriawinata, and S. Hassanpour, "Looking Under the Hood: Deep Neural Network Visualization to Interpret Whole-Slide Image Analysis Outcomes for Colorectal Polyps," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 7 2017, pp. 821–827.
- [72] P. Inglese and et al., "Deep learning and 3D-DESI imaging reveal the hidden metabolic heterogeneity of cancer," *Chemical Science*, pp. 3500–3511, 7 2017.
- [73] K. Zhang, X. Liu, F. Liu, L. He, L. Zhang, Y. Yang, W. Li, S. Wang, L. Liu, Z. Liu, X. Wu, and H. Lin, "An interpretable and expandable deep learning diagnostic system for multiple ocular diseases: Qualitative study," *Journal of Medical Internet Research*, vol. 20, no. 11, pp. 1–13, nov 2018. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/30429111> <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC6301833>
- [74] S. Zhang, Y. Luo, H. Du, Z. Jin, Y. Zhu, Y. Zhang, F. Xie, M. Zhang, X. Tian, and J. Zhang, "A Novel Interpretable Computer-Aided Diagnosis System of Thyroid Nodules on Ultrasound based on Clinical Experience," *IEEE Access*, pp. 1–1, 2020. [Online]. Available: <https://ieeexplore.ieee.org/document/9016204/>
- [75] Z. Zhang, P. Chen, M. McGough, F. Xing, C. Wang, M. Bui, Y. Xie, M. Sapkota, L. Cui, J. Dhillon, N. Ahmad, F. K. Khalil, S. I. Dickinson, X. Shi, F. Liu, H. Su, J. Cai, and L. Yang, "Pathologist-level interpretable whole-slide cancer diagnosis with deep learning," *Nature Machine Intelligence*, vol. 1, no. May, 2019.
- [76] X. Wang, Y. Peng, L. Lu, Z. Lu, and R. M. Summers, "TieNet: Text-Image Embedding Network for Common Thorax Disease Classification and Reporting in Chest X-Rays," *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pp. 9049–9058, 2018.
- [77] Z. Zhang, Y. Xie, F. Xing, M. McGough, and L. Yang, "MDNet: A Semantically and Visually Interpretable Medical Image Diagnosis Network," *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 3549–3557, 2017.
- [78] L. Shen, L. R. Margolies, J. H. Rothstein, E. Fluder, R. McBride, and W. Sieh, "Deep learning to improve breast cancer detection on screening mammography," *Scientific Reports*, vol. 9, no. 1, p. 12495, Aug 2019. [Online]. Available: <https://doi.org/10.1038/s41598-019-48995-4>
- [79] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, "Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization," in *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2017, pp. 618–626.
- [80] D. Bau, B. Zhou, A. Khosla, A. Oliva, and A. Torralba, "Network dissection: Quantifying interpretability of deep visual representations," *Proceedings - 30th IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017*, vol. 2017-Janua, pp. 3319–3327, 2017.
- [81] A. A. Kabbani, Y. Weerakkody, and et. al., "Breast imaging-reporting and data system (BI-RADS)," *Reston VA: American College of Radiology*, 1998.
- [82] B. Kim, M. Wattenberg, J. Gilmer, C. Cai, J. Wexler, F. Viegas, and R. Sayres, "Interpretability Beyond Feature Attribution: Quantitative Testing with Concept Activation Vectors (TCAV)," *ArXiv e-prints*, 2017.
- [83] S. T. Kim, J.-H. Lee, and Y. Ro, "Visual evidence for interpreting diagnostic decision of deep neural network in computer-aided diagnosis," in *Proceedings of Computer-Aided Diagnosis*, 03 2019, p. 19.
- [84] H2O, "H2O.ai," <https://www.h2o.ai/>, 2019, accessed: 2019-07-01.
- [85] A. Esteva, B. Kuprel, R. A. Novoa, J. Ko, S. M. Swetter, H. M. Blau, and S. Thrun, "Dermatologist-level classification of skin cancer with deep neural networks," *Nature*, vol. 542, no. 7639, pp. 115–118, Feb 2017.
- [86] B. Zhou, A. Khosla, À. Lapedriza, A. Oliva, and A. Torralba, "Learning Deep Features for Discriminative Localization," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 2921–2929.
- [87] Y. Luo, H.-H. Tseng, S. Cui, L. Wei, R. K. Ten Haken, and I. El Naqa, "Balancing accuracy and interpretability of machine learning approaches for radiation treatment outcomes modeling," *BJR—Open*, vol. 1, no. 1, p. 20190021, jul 2019. [Online]. Available: <https://www.birpublications.org/doi/10.1259/bjro.20190021>
- [88] J. T. Springenberg, A. Dosovitskiy, T. Brox, and M. Riedmiller, "Striving for Simplicity: The All Convolutional Net," *ArXiv e-prints*, 12 2014.
- [89] M. T. Ribeiro, S. Singh, and C. Guestrin, "Why Should I Trust You?": Explaining the Predictions of Any Classifier," in *Proceedings of the 22Nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, ser. KDD '16. New York, NY, USA: ACM, 2016, pp. 1135–1144.
- [90] K. Clark and et al., "The cancer imaging archive (TCIA): Maintaining and operating a public information repository," *Journal of Digital Imaging*, vol. 26, no. 6, pp. 1045–1057, 2013.
- [91] TCGA, "The cancer genome atlas," <https://www.cancer.gov/tcga>, 2019, accessed: 2019-07-01.
- [92] L. Van Der Maaten and G. Hinton, "Visualizing Data using t-SNE," *Journal of Machine Learning Research*, vol. 9, pp. 2579–2605, 2008.
- [93] M. Du, N. Liu, and X. Hu, "Techniques for interpretable machine learning," *ArXiv*, vol. abs/1808.00033, 2019.
- [94] N. Carlini, C. Liu, Ú. Erlingsson, J. Kos, and D. Song, "The secret sharer: Evaluating and testing unintended memorization in neural networks," in *28th USENIX Security Symposium (USENIX Security 19)*. Santa Clara, CA: USENIX Association, Aug. 2019, pp. 267–284. [Online]. Available: <https://www.usenix.org/conference/usenixsecurity19/presentation/carlini>
- [95] S. G. Finlayson, J. D. Bowers, J. Ito, J. L. Zittrain, A. L. Beam, and I. S. Kohane, "Adversarial attacks on medical machine learning," *Science*, vol. 363, no. 6433, pp. 1287–1289, 2019. [Online]. Available: <https://www.science.org/doi/abs/10.1126/science.aaw4399>
- [96] C. J. Kelly, A. Karthikesalingam, M. Suleyman, G. Corrado, and D. King, "Key challenges for delivering clinical impact with artificial intelligence," *BMC Medicine*, vol. 17, no. 1, pp. 1–9, oct 2019. [Online]. Available: <https://bmcmmedicine.biomedcentral.com/articles/10.1186/s12916-019-1426-2>
- [97] J. R. Zech, M. A. Badgeley, M. Liu, A. B. Costa, J. J. Titano, and E. K. Oermann, "Variable generalization performance of a deep learning model to detect pneumonia in chest radiographs: A cross-sectional study," *PLOS Medicine*, vol. 15, no. 11, pp. 1–17, 11 2018.
- [98] M. Havaei, A. Davy, D. Warde-Farley, A. Biard, A. Courville, Y. Bengio, C. Pal, P.-M. Jodoin, and H. Larochelle, "Brain tumor segmentation with deep neural networks," *Medical Image Analysis*, vol. 35, p. 18–31, Jan 2017. [Online]. Available: <http://dx.doi.org/10.1016/j.media.2016.05.004>
- [99] M. Mittal, L. M. Goyal, S. Kaur, I. Kaur, A. Verma, and D. J. Hemanth, "Deep learning based enhanced tumor segmentation approach for mr brain images," *Appl. Soft Comput.*, vol. 78, pp. 346–354, 2019.
- [100] H. Kim, J. Jung, J. Kim, B. Cho, J. Kwak, J. Y. Jang, S.-w. Lee, J.-G. Lee, and S. M. Yoon, "Abdominal multi-organ auto-segmentation using 3d-patch-based deep convolutional neural network," *Scientific Reports*, vol. 10, no. 1, p. 6204, Apr 2020. [Online]. Available: <https://doi.org/10.1038/s41598-020-63285-0>
- [101] W. Zhu, L. Xie, J. Han, and X. Guo, "The application of deep learning in cancer prognosis prediction," *Cancers*, vol. 12, no. 3, p. 603, Mar 2020.
- [102] Y. Xu, A. Hosny, R. Zeleznik, C. Parmar, T. Coroller, I. Franco, R. H. Mak, and H. J. Aerts, "Deep learning predicts lung cancer treatment response from serial medical imaging," *Clinical Cancer Research*, vol. 25, no. 11, pp. 3266–3275, 2019.
- [103] F. Liu, P. P. Yadav, A. M. Baschnagel, and A. B. McMillan, "Mr-based treatment planning in radiation therapy using a deep learning approach," *Journal of Applied Clinical Medical Physics*, vol. 20, pp. 105 – 114, 2019.
- [104] Y. Fu, Y. Lei, T. Wang, W. J. Curran, T. Liu, and X. Yang, "Deep learning in medical image registration: A review," *ArXiv e-prints*, 2019.

[105] H. Chen, Y. Zhang, M. K. Kalra, F. Lin, Y. Chen, P. Liao, J. Zhou, and G. Wang, "Low-dose ct with a residual encoder-decoder convolutional neural network," *IEEE Transactions on Medical Imaging*, vol. 36, no. 12, p. 2524–2535, Dec 2017. [Online]. Available: <http://dx.doi.org/10.1109/TMI.2017.2715284>

[106] Q. Yang, P. Yan, Y. Zhang, H. Yu, Y. Shi, X. Mou, M. K. Kalra, Y. Zhang, L. Sun, and G. Wang, "Low-dose ct image denoising using a generative adversarial network with wasserstein distance and perceptual loss," *IEEE Transactions on Medical Imaging*, vol. 37, no. 6, p. 1348–1357, Jun 2018. [Online]. Available: <http://dx.doi.org/10.1109/TMI.2018.2827462>

[107] X. Zhang, N. Wang, H. Shen, S. Ji, X. Luo, and T. Wang, "Interpretable Deep Learning under Fire," *ArXiv e-prints*, 2018. [Online]. Available: <http://arxiv.org/abs/1812.00891>

[108] A. Adegun and S. Viriri, "Deep learning model for skin lesion segmentation: Fully convolutional network," in *Image Analysis and Recognition*, F. Karray, A. Campilho, and A. Yu, Eds. Cham: Springer International Publishing, 2019, pp. 232–242.

[109] J. Lameski, A. Jovanov, E. Zdravevski, P. Lameski, and S. Gievaska, "Skin lesion segmentation with deep learning," in *IEEE EUROCON 2019 -18th International Conference on Smart Technologies*, 2019, pp. 1–5.

[110] H.-M. Zhang and B. Dong, "A review on deep learning in medical image reconstruction," *Journal of the Operations Research Society of China*, Jan 2020. [Online]. Available: <https://doi.org/10.1007/s40305-019-00287-4>

[111] P.-J. Kindermans, S. Hooker, J. Adebayo, M. Alber, K. T. Schütt, S. Dähne, D. Erhan, and B. Kim, *The (Un)reliability of Saliency Methods*. Springer International Publishing, 2019, ch. 4, pp. 267–280.

[112] J. Adebayo, J. Gilmer, M. Muelly, I. Goodfellow, M. Hardt, and B. Kim, "Sanity checks for saliency maps," in *Advances in Neural Information Processing Systems 31*. Curran Associates, Inc., 2018, pp. 9505–9515.

[113] J. Fauw, J. Ledsam, B. Romera-Paredes, S. Nikolov, N. Tomasev, S. Blackwell, H. Askham, X. Glorot, B. O'Donoghue, D. Visentin, G. Driessche, B. Lakshminarayanan, C. Meyer, F. Mackinder, S. Bouton, K. Ayoub, R. Chopra, D. King, A. Karthikesalingam, and O. Ronneberger, "Clinically applicable deep learning for diagnosis and referral in retinal disease," *Nature Medicine*, vol. 24, 09 2018.

[114] D. Ravi, C. Wong, F. Deligianni, M. Berthelot, J. Andreu, B. Lo, and G.-Z. Yang, "Deep learning for health informatics," *IEEE journal of biomedical and health informatics*, vol. PP, 12 2016.



José P. Amorim José Pereira Amorim is a PhD student at the University of Coimbra. He got his M.S. degree in Informatics Engineering at Faculty of Engineering of the University of Porto in 2017. He is currently affiliated with the IPO-Porto Research Centre (CI-IPOP) and the Centre for Informatics and Systems of the University of Coimbra (CISUC) and his research is mainly focused on creating interpretable models in order to assist physicians in the field of oncology.




Pedro H. Abreu Pedro Henriques Abreu is an Assistant Professor at DEI, University of Coimbra, Portugal. He got his Informatics Engineering Degree from Porto University in 2006 and a PhD in Soccer Teams Modeling from the Porto University (2011). His interests include (but are not limited to) Artificial Intelligence applied to Medical Informatics contexts. Pedro Henriques Abreu is the author of more than sixty publications in international conferences and journals.



Alberto Fernández Alberto Fernández received the M.Sc. and Ph.D. degrees in computer science from the University of Granada, Spain, in 2005 and 2010, respectively. He has been a University Professor in the Department of Computer Science and Artificial Intelligence since 2017 and teaches at the Faculty of Science and at the E.T.S. of Computer Engineering and Telecommunications at the University of Granada, Spain. He has received the Extraordinary Doctoral Thesis Award 2010, on three occasions the Award for Works of Excellence of the University of Granada, as well as the Lofti Zadeh award for best publication in 2009. His research focuses on topics such as classification in unbalanced domains, learning fuzzy rules, evolutionary algorithms, multi-classification problems with Ensembles and decomposition techniques and Data Science and Big Data applications, in which he has published more than 50 works in international journals ISI, collaborated in more than 20 projects and research contracts and 3 doctoral theses. He is a member of the editorial committee of a number of international journals such as Applied Intelligence, PLOS-One, Progress in Artificial Intelligence, Cognitive Computation and Big Data, and a member of the "referees" committee of other ISI journals such as Fuzzy Sets and Systems, IEEE Transactions on Fuzzy Systems, Information Sciences, among others.



Mauricio Reyes Mauricio Reyes is currently the head of the Healthcare Imaging A.I. group at the University of Bern, Switzerland, as well as head of the Data Science team of the Insel Data Science Center (IDSC) of the University Hospital Bern. Before that he received his bachelor's degree at the University of Santiago de Chile, Chile in 2001. His thesis Three-dimensional Reconstruction of a Human Embryo Hand Using Artificial Vision Techniques was awarded Best Electrical Engineering bachelor thesis work. Later on, during 2002- 2004, he conducted studies to obtain a Ph.D. degree in computer sciences from the University of Nice, France on the topic of lung cancer imaging and breathing compensation in emission tomography, under the supervision of Prof. Dr. Grégoire Malandain, Asclepius research project (formerly known as Epi-daure). In 2006 he joined the Medical Image Analysis group at the MEM Research Center as a postdoctoral fellow focusing on topics related to medical image analysis and statistical shape models for orthopaedic research. In 2008, he took over the lead of the Medical Image Analysis group at the Institute for Surgical Technology and Biomechanics, Switzerland.



João Santos João Santos is a physicist with a PhD in Physics (Condensed Matter) and a MSc in Physics (Optoelectronics and Lasers). He concluded his Residency in Medical Physics in 2005 at the Portuguese Institute of Oncology of Porto (IPOPFG, EPE). He is currently working as a Medical Physicist Expert at IPOPFG, EPE and he is the Coordinator of the Medical Physics, Radiobiology and Radiation Protection group of the IPOPFG, EPE Research Center since 2008.



Miguel H. Abreu Miguel Henriques Abreu, M.D is an oncologist in Instituto Português Oncologia do Porto Francisco Gentil. Graduated in Medicine, University of Porto in 2007 and PhD in Medical Sciences from the same Institution in 2016. He completed his residence training in Medical Oncology at Portuguese Institute of Oncology of Porto in 2013 and become medical assistant in the same Institution since 2014, being dedicated to the treatment of breast and gynaecological cancer patients. He is an active member of important research groups including the

EORTC group for gynaecological cancer, EORTC for breast cancer and the SOLTI group, and is coinvestigator in 37 clinical trials and principal investigator in 4 of them.