CIDAI Centre of Innovation for Data tech and Artificial Intelligence

T3 EURECAT July 2022

Health medical image projects

Explainable and interpretable AI

Vicent Ribas Ripoll
Felip Miralles Barrachina
Xavier Rafael-Palou
Paula Subías Beltrán

CIDAI-PAI 01-2021-D03

Some rights reserved. This work is available under the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International (CC BY-NC-SA 4.0).

Under the terms of this licence, you may copy, redistribute and adapt the work for non-commercial purposes, provided the work is appropriately cited, as indicated below.
In any use of this work, there should be no suggestion that CIDAI endorses any specific organization, products or services. The use of the CIDAI logo is not permitted.
If you adapt the work, then you must license your work under the same or equivalent Creative Commons licence.

If you create a translation of this work, you should add the following disclaimer along with the suggested citation: "This translation was not created by the Centre of Innovation for Data tech and Artificial Intelligence (CIDAI). CIDAI is not responsible for the content or accuracy of this translation. The original English edition shall be the binding and authentic edition".
Any mediation relating to disputes arising under the licence shall be conducted in accordance with the mediation rules of the World Intellectual Property Organization.

Suggested citation. CIDAI-PAI 01-2021-D03// Health medical image projects - T3. Explainable and interpretable AI by Vicent Ribas Ripoll, Felip Miralles Barrachina, Xavier Rafael-Palou and Paula Subías Beltrán, CIDAI, EURECAT 2021. Licence: CC BY-NC-SA 4.0

1. INTRODUCTION 4
2. BACKGROUND 6
3. METHODS 9
4. RESULTS 11
5. CONCLUSIONS & FUTURE WORK 16
6. CITED WORK 18

Introduction

1

Artificial intelligence (AI) has achieved growing momentum in its application in healthcare to deal with the increased complexity, scalability, and automation, and that also permeates day-to-day clinical practice. A rapid surge in the complexity and sophistication of AI-powered systems has evolved to such an extent that humans do not understand the complex mechanisms by which AI systems work or how they make certain decisions — something that is particularly a challenge when AI-based systems compute outputs that are unexpected or seemingly unpredictable. This especially holds true for opaque decision- making systems, such as those using deep neural networks (DNNs), which are considered complex black box models. The inability for humans to see inside black boxes can result in AI adoption (and even its further development) being hindered, which is why growing levels of autonomy, complexity, and ambiguity in AI methods continues to increase the need for interpretability, transparency, understandability, and explainability of AI products/outputs (such as predictions, decisions, actions, and recommendations). Not only are these elements crucial to ensuring that humans can understand and, consequently, trust AI systems but they are fundamental for regulatory compliance (mainly with the new GDPR) of AI-based clinical products. GDPR enables patients to ask for a second opinion about any decision made about their health regardless of it being made by a human or a computer. Explainable artificial intelligence (XAI) refers to methods and techniques that produce accurate, explainable models of why and how an AI algorithm arrives at a specific decision so that AI solution results can be understood by humans.

Without explanations behind an AI model's internal functionalities and the decisions it makes, there is a risk that the model would not be considered trustworthy or legitimate. XAI provides the needed understandability and transparency to enable greater trust toward AI- based solutions. Thus, XAI is acknowledged as a crucial feature for the practical deployment of AI models in systems and, more importantly, for satisfying the fundamental rights of AI users related to AI decision-making (according to European Commission ethical guidelines for trustworthy AI). Standardization bodies such as the European Telecommunications Standards Institute (ETSI) and the Institute of Electrical and Electronics Engineers Standards Association (IEEE SA) also emphasize the importance of XAI where AI models are deployed, indicating XAI's growing importance in the future [EUT1]. AI deployers and developers must comply with these ethical guidelines and regulations to ensure that their AI solutions are explainable and trustable [EUT2].

The aim of this short deliverable is to present the work on XAI for the different use cases selected for the PAI – CIDAI.

1.1. Case-1

Up4Health-Characterization of Lung Pathologies

1.2. Case-2

Nodule malignancy assessment through convolutional neural networks (Deep Lung).

1.3. Case-3

AI for Kidney anomaly detection

Background

2

Building trust is essential for users to accept AI-based solutions and those systems incorporating decisions made by them. There are, however, significant challenges in developing explainability methods. One of them is the trade-off between attaining the simplicity of algorithm transparency and impacting the high-performing nature of complex but opaque models (when one increases the transparency aspect, privacy and the security of sensitive data come into question). Yet another challenge is to identify the right information for the user, where different levels of knowledge will come into play. Beyond selecting the level of knowledge retained by the user, generating a concise (simple but meaningful) explanation also becomes a challenge. Most explainability methods focus on explaining the processes behind an AI decision, which is sometimes agnostic to the context of its application, providing unrealistic explanations. Researchers attempt to integrate knowledge-based systems so that the explanation becomes relevant to its application's context.
XAI helps deliver trust by supporting with the following properties:

- Trustworthiness, to attain the trust of humans on the AI model by explaining the characteristic and rationale of the AI output
- transferability, where the explanation of an AI model allows better understanding of it so that it can be transferred to another problem or domain/application properly
- informativeness, relating to informing a user regarding how an AI model works to avoid misconception (this is also related to human agency and autonomy, which ensures humans understand AI outcomes and can take intervening actions on that basis)
- confidence, which is achieved through having a model that is robust, stable, and explainable to support human confidence in deploying an AI model
- privacy awareness, ensuring that the AI and XAI methods do not expose private data (which can be done through data anonymization)
- actionability, with XAI providing indications regarding how a user could change an action to yield a different outcome in addition to providing the rationale for an outcome
- tailored (user-focused) explanations, allowing humans — as AI system users of different knowledge backgrounds — to understand the behaviour and predictions made by AI-based systems through tailored explanations based on their roles, goals, and preferences.

2.1. CVC Up4Health-Characterization of Lung Pathologies

Many works addressing radiological interpretation of deep networks rely on the identification of the image regions that the network pays attention to. This identification is computed by back-propagating the content of intermediate layers and using attention models [CVC2]. Attention models have gained increasing interest to improve segmentation, detection, and final diagnosis classification of lesions [CVC3, CVC4]. From these works, it has been demonstrated that the use of self-attention in a transformer architecture accurately improves results [CVC5].

Given that the attention mechanism dynamically computes a weight vector to extract partial image features supporting the final prediction, the visualization of the associated attention maps allows the identification of regions relevant to the model's output [CVC6]. Visualization of attention maps provides a level of transparency that allows users to validate model outputs, but still lacks the link to the clinical domain knowledge. A recent study [CVC7] shows that to facilitate the use of AI models in clinical practice, clinicians need to justify their decision-making based on the model's prediction in terms of the current evidence based medical practice. Indeed, the combination of radiological and visual deep features related to malignancy has shown to improve interpretability of results [CVC8]. Therefore, some recent works addressing automatic report generation have been published. In [CVC9, CVC10], the authors present models that combine attention maps and radiological reports to predict lesion pathology and automatically generate reports for new cases. Also, preliminary studies with simple forms of self-attention have shown its usefulness in medical imaging segmentation, classification, detection, reconstruction, registration, synthesis, and clinical report generation [CVC11].

Finally, in [CVC12] several models directly linking abstract features to radiological descriptions are presented and tested for the characterization and diagnosis of lesions. All these works are end-to-end approaches that use a combination of medical images and radiological annotations. Thus, they are unable to provide an interpretation of pre-trained visual features.

2.2. Eurecat Deep Lung – Nodule Malignancy Classification

An ML model is built using a certain method based on the properties of the available data. A state-of-the-art model, such as Deep Lung, is used to achieve very high performance in predicting or classifying instances; however, as the complexity of the models is continuously increasing, it is becoming harder to understand how particular predictions are made and which factors contributed (and how much) to a given prediction. XAI methods are used to capture and explain the characteristics of complex ML models. When a prediction is complemented with an explanation, it is expected that the ML model will gain more trust from humans.

Explainability for ML models requires the data, an ML model, and the output/prediction of the ML model as input to generate the explanation. The explanations are presented directly to humans or analyzed further by a Machine Reasoning (MR) system. Having not only prediction but also explanation from ML models enriches information for MR in making decisions (for example, in fulfilling an intent). Without ML explainability, one could not be sure whether the model complies with the rules or standards that the model should fulfill (for example, being not biased towards sensitive features).

Explainability for ML can be used at different stages in developing an AI system. Once a black box ML model is built with satisfactory performance, XAI methods (for example, SHAP [EUT3], XGBoost [EUT4], Causal Dataframe [EUT5], PI [EUT6], and so on) are applied to obtain the general behavior of a model (also known as “global explanation”). While the global explanation, which explains the whole data, is important, it is also necessary to obtain the explanation of every single data prediction (also known as “local explanation”). Whenever new/more data is fed into the model, the rationale of every ML prediction is presented in a local manner, where one explanation may be different from another. Some methods mostly used for local explainability include LIME [EUT7], ELI5 [EUT8], CFProto [EUT9], and more. For the case of Deep Lung, we have chosen the Grad-CAM pipeline by Selejvarju et al [EUT10].

2.3. NTTDATA - AI for Kidney anomaly detection

The AI for kidney anomalies detection project is part of a larger initiative which has allocated specific resources in the scope of XAI. Given the project timeline, the scope of the tasks proposed has not included to deal with XAI as it is being worked out in the scope of the larger project. Shall it be the case any specific insights can be derived from that work it will be included.

Deployment and execution of AI models. Fine tuning and adjustments

3.1. CVC Up4Health-Characterization of Lung Pathologies

Current deep approaches are black-box models that provide clinical outcomes directly from scans and, thus, their output is even harder to interpret than less accurate systems based on hand-crafted features explicitly escribing the morphology and shape of lesions. Although activation and attention maps improve interpretability and trustworthiness of deep approaches, they do not have a straight-forward correspondence with the annotations that a radiologist would made to summarize the clinical findings of the lesion.

We formulate interpretability of models as a transformation between the domains of radiomic visual features and radiological descriptions. We propose to use transformers to map abstract radiomic features to specific radiological annotations. In the radiomic domain, lesions are described as a sentence given by a self-defined embedding of their visual features. This sequence of visual tokens is the input to a multi-attention transformer optimizing a hybrid loss that translates the visual sequence into a sequence of radiological words describing the lesion appearance in quantitative and qualitative terms. The network hyperparameters are optimized using a Non-dominated Sorting Genetic Algorithm. Results obtained for 584 patients of the LIDC-IDRI [CVC13] public data base outperform state-of-art methods in the prediction of radiological annotations.

3.2. Eurecat Deep Lung – Nodule Malignancy Classification

The Grad-CAM method proposes a technique for producing “visual explanations” for decisions from a large class of CNN-based models, making them more transparent. The - Gradient-weighted Class Activation Mapping (Grad-CAM) approach, uses the gradients of any target concept, flowing into the final convolutional layer to produce a coarse localization map highlighting important regions in the image for predicting the concept. Grad-CAM is applicable to a wide variety of CNN model-families: (1) CNNs with fully connected layers, (2) CNNs used for structured outputs, (3) CNNs used in tasks with multimodal inputs or reinforcement learning, without any architectural changes or re-training. We combine Grad-CAM with fine-grained visualizations to create a high-resolution class-discriminative visualization and apply it the models developed in Deep Lung including ResNet-based architectures. In the context of image classification models, our visualizations (a) lend insights into their failure modes, (b) are robust to adversarial images, (c) outperform previous methods on localization, (d) are more faithful to the underlying model and (e) help achieve generalization by identifying dataset bias.

Model Validation

4

4.1. CVC-Up4Health- Characterization of lung pathologies

Two different experiments were carried out:

1. Model Optimization. A training and selection of the best hyperparameters.
2. Model Verification. A validation of the best model on an independent set of patients to assess the reproducibility of results including a comparison to [CVC11].

In these experiments we have used the LIDC-IDRI [CVC13] public data base, which contains 1018 CT scans, each annotated by four radiologists and comprising nodule location, segmentation, characterization, and malignancy. From this, a subset of 584 patients with 1110 nodules corresponding to the LUNA16 challenge [CVC14] were used.

The radiological features are subtlety, internal structure, calcification, sphericity, margin, lobulation, spiculation and texture as well as malignancy. Calcification, internal structure, sphericity and subtlety are qualitative descriptors and, thus, are considered categorical attributes for the definition of the joint loss. The remaining ones, margin, lobulation, spiculation and texture, are degrees of the pathological structure and are considered quantitative attributes. The multiple radiological annotations are summarized as the most frequent label for categorical attributes and the average for quantitative ones. The input visual features are 19 GLCM Pyradiomics [CVC15] (version 3.01) features selected according to their correlation to lesion malignancy as described in [CVC16]. Texture features were computed in 2D slices, which provides a total number of 6751 samples.

Model Optimization. The accuracy of the best model has been assessed using a cross-entropy accuracy for the categorical variables and mean squared error for the quantitative ones. The ranges obtained for the outer folds are, respectively, 1e −3 ±1e −4 for classification accuracy and 1.05 ±0.27 for quantification error. The best hyperparameters' values found in the inner CV are presented in Table 1. These values were used to train the models in the outer CV which losses convergence are shown in the Figure 1.

Table 1. Hyperparameter search space for our model architecture

Hyper-parameter	nhead	nhid	emsize	dropout	nlayers	lr	alpha
Search space	[2, 10]	[8, 300]	[2,15]	[0.2, 0.5]	[1, 3]	[e-4, e-1]	[0.1, 0.9]
Best values	2	227	11	0.2	3	0.06	0.83

Figure 1. Convergenceof the outer CV folds.

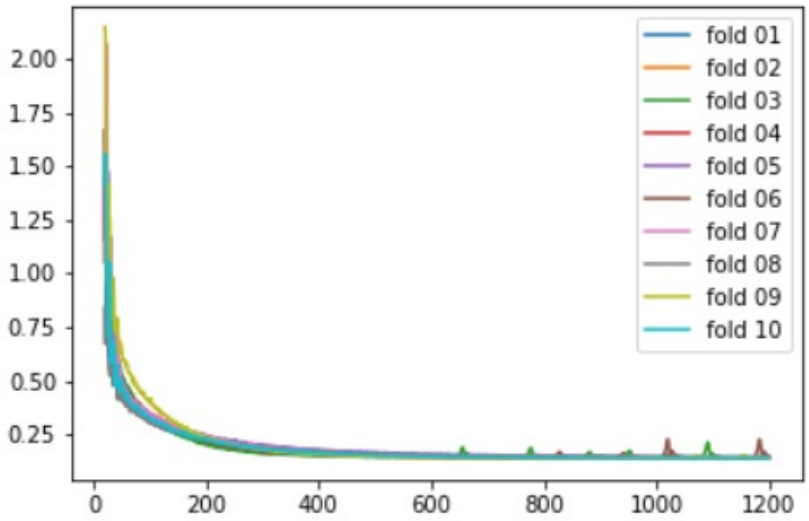

Model Verification. The performance of our best model has been verified in terms of accuracy in the classification of categorical variables and mean square error for the quantitative ones. The average classification accuracy for the hold-out set is 0.0005, while the average quantitative error is 1.51. Our model has been also compared to the accuracy of the feature predictions reported in [CVC12]. To do so, for each attribute the original LIDC-IDRI ordinal labels have been encoded into binary classes, corresponding to the intervals defined in [CVC12]. The average accuracy of the feature predictions of each model is shown in Table 2 with best performance in boldface. Our method has average prediction of 84.57 ±16.12 in comparison to the average 76.28 ±8.44 achieved by [CVC12]. Discarding sphericity, average predictions for the remaining features are significantly higher (pvalue=0.007). The drop-in detection of sphericity (which has the lower prediction for both methods) can be attributed to the fact that being a 3D shape attribute it could be prone to have larger error in visual inspections.

Table 2. Comparison to the state-of-art in terms of accuracy classification in radiological features.

	Calcification	Lobulation	Margin	Sphericity	Spiculation	Subtlety	Texture
Our method	0.96	0.90	0.85	0.46	0.88	0.93	0.94
[CVC12]	0.95	0.75	0.76	0.66	0.75	0.70	0.77

4.2. Eurecat Deep Lung – Nodule Malignancy Classification

Three different methods (3DCNN-MAL, TS-3DCNN and TS-3DCNN-MAL) were evaluated on the resulting 32 matching lung nodules (65% of them cancerous) from the re-identification step. For the evaluation of this methods, we directly used the 3DCNN-MAL classifier on the evaluation cases, while for the other two classifiers, to avoid data leakage for this evaluation, we retrained them before being evaluated, using the training partition of the VHLung with a 10-fold cross-validation. Particularly for the 3DCNN-MAL classifier, as it outputs 3 probabilities, we assumed cancer prediction when this classifier reported as maximum probability either the category suspicious or malignancy.

Table-4 shows the resulting classification performances for these models on the 32 matching nodules. We computed, precision (PREC), recall (REC) and specificity (SPEC), as well as balanced accuracy (BA). Due to the data was unbalanced towards the cancer case, the BA was used as the reference metric. Reported values in Table-4 are the mean and 2 standard errors (associated with a 95% of confidence).

Table-4. Performance of the different nodule malignancy classifiers of component of the pipeline.

	TEST			
	BA	PREC	REC	SPEC
3DCNN-MAL	0.776+/-0.153	0.808+/-0.152	1.0+/-0.0	0.552+/-0.307
TS-3DCNN	0.810+/-0.203	0.899+/-0.167	0.791+/-0.309	0.829+/-0.296
TS-3DCNN-MAL	0.825+/-0.201	0.910+/-0.179	0.821+/-0.33	0.830+/-0.348

The 3DCNN-MAL classifier obtained a balanced accuracy score of 0.77, while the TS-3DCNN achieved a 0.81. However, the TS-3DCNN-MAL, which integrated the outcomes of the 3DCNN-MAL model, improved the balanced accuracy score of the TS-3DCNN model, a 1.5%. For further comparison of these models, we show Figure-9 with the ROC-curves of the two best models.

Figure 9. Roc curves of the nodule malignancy classification models

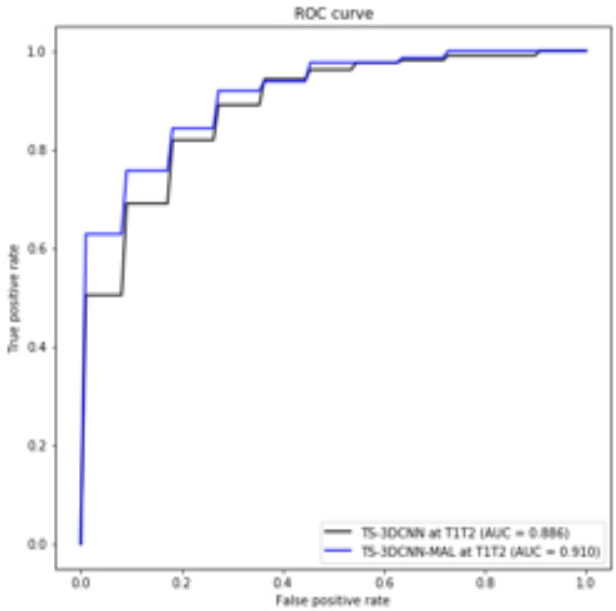

For a further intuition and visual interpretability of the areas of the images that the TS-3DCNN-MAL network took most seriously in deciding which class to assign to the image, we extracted the Grad-CAMs features for this classifier. On the last layer of the first block of the TS-3DCNN-MAL network we obtained the gradients and the feature activations, and they were multiplied after being pooled on the channel dimension. Figures-10,11 show the results of this visualization technique on the three-planes of different lung nodules (either benign or malign).

Figure-10. Grad-CAM features from the TS-3DCNN-TL-MAL network for 2 malign nodules

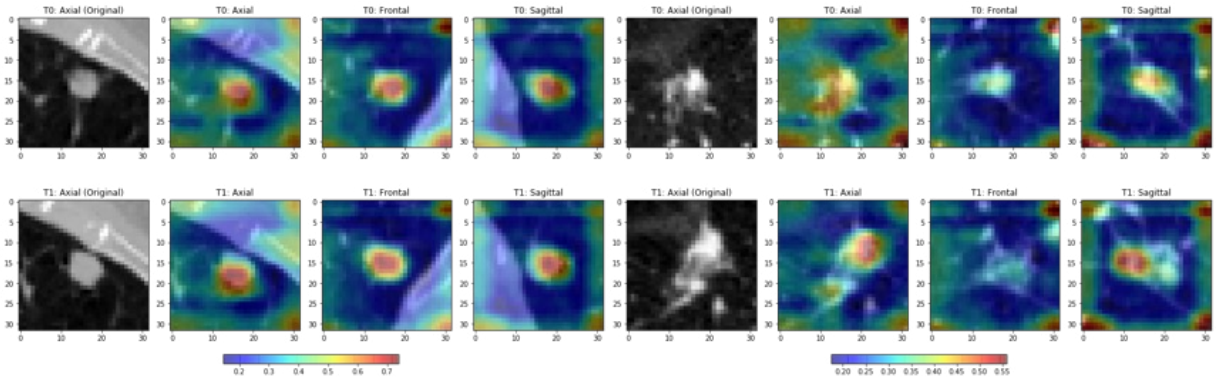

Figure-11. Grad-CAM features from the TS-3DCNN-TL-MAL network for 2 benign nodules

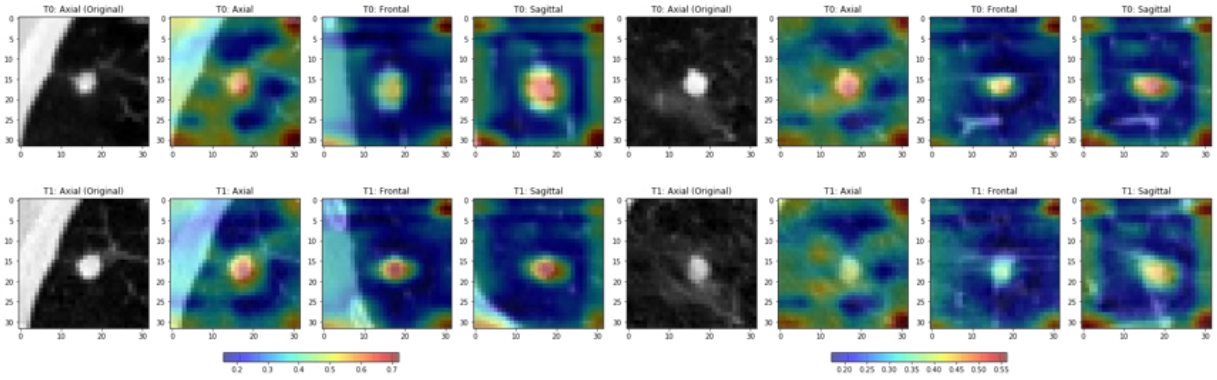

Conclusions & future work

5

In this short document we presented the 'explainable AI' (XAI) activities developed by two use cases of elected for the PAI – CIDAI. More particularly, the use case Up4Health-Characterization of Lung Pathologies proposed an interpretability model as a transformation between the domains of radiomic visual features and radiological description using transformers. These transformers map abstract radiomic features to specific radiological annotations. In the radiomic domain, lesions are described as a sentence given by a self-defined embedding of their visual features. This sequence of visual tokens is the input to a multi-attention transformer optimizing a hybrid loss that translates the visual sequence into a sequence of radiological words describing the lesion appearance in quantitative and qualitative terms. The radiological features considered in this use case are subtlety, internal structure, calcification, sphericity, margin, lobulation, spiculation and texture as well as malignancy. Calcification, internal structure, sphericity, and subtlety are qualitative descriptors and, thus, are considered categorical attributes for the definition of the joint loss. The remaining attributes (margin, lobulation, spiculation and texture) are degrees of the pathological structure and are considered quantitative attributes. The multiple radiological annotations are summarized as the most frequent label for categorical attributes and the average for quantitative ones. This method presents an average accuracy of the feature predictions of each model of 84.57 ±16.12.

The use-case Deep Lung uses the Grad-CAM method for producing 'visual explanations' of the decisions from a large class of CNN-based models. This approach uses the gradients of any target concept flowing into the final convolutional layer to produce a coarse localization map highlighting the most important regions in the image for prediction. Deep Lung combines Grad-CAM with fine-grained visualizations to create high-resolution class-discriminative visualizations that are applied in the different models evaluated in the project. Different Grad-CAM examples have been presented for both benign and malignant lung nodules.

Cited work

6

[EUT1] Frost, L., Meriem, T. B., Bonifacio, J. M., Cadzow, S., Silva, F. d., Essa, M., Forbes, R., Marchese, Odini, M., Sprecher, N., Toche, C., & Wood, S. (2020). Artificial Intelligence and future directions for ETSI. ETSI White Paper No. #34, 1 (ISBN 979-10-92620-30- 1).

[EUT2] Anneroth, (2019). Responsible AI – a human right?

[EUT3] Lundberg, S., & Lee, S.-I. (2017). A Unified Approach to Interpreting Model Predictions. Advances in Neural Information Processing Systems 30 (NIPS 2017). Curran Associates, Inc. pp. 4765-4774.

[EUT4] Chen, T., & Guestrin, C. (2016). XGBoost: A Scalable Tree Boosting System. Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. San Francisco, California, USA. Association for Computing Machinery. pp. 785- 794.

[EUT5] Kelleher, A. (2017). Causality.

[EUT6] Altmann, A., Toloşi, L., Sander, O., & Lengauer, T. (2010). Permutation importance: a corrected feature importance measure. Bioinformatics, 10, 1340-1347.

[EUT7] Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). "Why Should I Trust You?": Explaining the Predictions of Any Classifier. 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. San Francisco, California, USA. Association for Computing Machinery.

[EUT8] Korobov, , & Lopuhin, K. (2019). ELI5 Documentation, Release 0.9.0.,

[EUT9] Looveren, A. V., & Klaise, J. (2019). Interpretable Counterfactual Explanations Guided by Prototypes. Arxiv. [Preprint]

[EUT10] The interpretability methods used in Deep Lung are based on the work of Selvaraju R., Ramprasaath R., et al 'Grad-cam: Visual explanations from deep networks via gradient-based localization', Proc. IEEE International Conference on Computer Vision, 2017.

[CVC1] Ramírez, E., Sanchez, C., & Gil, D. (2021). BronchoX-Cloud: A Diagnostic as a Service Cloud Platform for Planning and Navigation Bronchoscopy. In 2021 23rd International Symposium on Symbolic and Numeric Algorithms for Scientific Computing (SYNASC) (pp. 227-233). IEEE.

[CVC2] Vasilakakis, M., Sovatzidi, G., Iakovidis, D.K.: Explainable classification of weakly annotated wireless capsule endoscopy images based on a fuzzy bag-of-colour features model and brain storm optimization. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 488–498. Springer (2021)

[CVC3] Gao, Y., Zhou, M., Metaxas, D.N.: Utnet: a hybrid transformer architecture for medical image segmentation. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 61–71. Springer (2021)

[CVC4] Jiang, H., Shen, F., Gao, F., Han, W.: Learning efficient, explainable and discriminative representations for pulmonary nodules classification. Pattern Recognition 113, 107825 (2021)

[CVC5] Shamshad, F., Khan, S., Zamir, S.W., Khan, M.H., Hayat, M., Khan, F.S., Fu, H.: Transformers in medical imaging: A survey. arXiv preprint arXiv:2201.09873 (2022).

[CVC6] Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929 (2020)

[CVC7] Tonekaboni, S., Joshi, S., McCradden, M.D., Goldenberg, A.: What clinicians want: contextualizing explainable machine learning for clinical end use. In: Machine learning for healthcare conference. pp. 359–380. PMLR (2019)

[CVC8] Zheng, S., Shen, Z., Pei, C., Ding, W., Lin, H., Zheng, J., Pan, L., Zheng, B., Huang, L.: Interpretative computer-aided lung cancer diagnosis: From radiology analysis to malignancy evaluation. Computer Methods and Programs in Biomedicine 210, 106363 (2021)

[CVC9] Lee, H., Kim, S.T., Ro, Y.M.: Generation of multimodal justification using visual
word constraint model for explainable computer-aided diagnosis. In: Interpretability of Machine Intelligence in Medical Image Computing and Multimodal Learning for Clinical Decision Support, pp. 21–29. Springer (2019)

[CVC10] Wang, X., Peng, Y., Lu, L., Lu, Z., Summers, R.M.: Tienet: Text-image embedding network for common thorax disease classification and reporting in chest x-rays. In:Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 9049–9058 (2018)

[CVC11] Shamshad, F., Khan, S., Zamir, S.W., Khan, M.H., Hayat, M., Khan, F.S., Fu, H.: Transformers in medical imaging: A survey. arXiv preprint arXiv:2201.09873 (2022)

[CVC12] Marques, S., Schiavo, F., Ferreira, C.A., Pedrosa, J., Cunha, A., Campilho, A.: A multi-task cnn approach for lung nodule malignancy classification and characterization. Expert Systems with Applications 184, 115469 (2021)

[CVC13] Armato III, S.G., McLennan, G., Bidaut, L., McNitt-Gray, M.F., Meyer, C.R., Reeves, A.P., Zhao, B., Aberle, D.R., Henschke, C.I., Hoffman, E.A., et al.: The lung image database consortium (lidc) and image database resource initiative (idri): a completed reference database of lung nodules on ct scans. Medical physics 38(2), 915–931 (2011)

[CVC14] Setio, A.A.A., Traverso, A., De Bel, T., Berens, M.S., Van Den Bogaard, C., Cerello, P., Chen, H., Dou, Q., Fantacci, M.E., Geurts, B., et al.: Validation, comparison, and combination of algorithms for automatic detection of pulmonary noddules in computed tomography images: the luna16 challenge. Medical image analysis
42, 1–13 (2017)

[CVC15] van Griethuysen, J.J., Fedorov, A., Parmar, C., Hosny, A., Aucoin, N., Narayan, V., Beets-Tan, R.G., Fillion-Robin, J.C., Pieper, S., Aerts, H.J.: Computational radiomics system to decode the radiographic phenotype. Cancer research 77(21), e104–e107 (2017).

[CVC16] Torres, G., Baeza, S., Sanchez, C., Arriaga, I., Rosell, A., Gil, D.: An intelligent radiomic approach for lung cancer screening. Applied Sciences (2022).

[CVC17]

CIDAI Centre of Innovation for Data tech and Artificial Intelligence

C/Bilbao, 72 Edif. A.
08005 - Barcelona
Tel. +34 93 7419 100
info@cidai-catalonia-ai.eu
www.cidai.eu