

Article

Evaluating Explainability: A Framework for Systematic Assessment of Explainable AI Features in Medical Imaging

Miguel A. Lago *, Ghada Zamzmi , Brandon Eich and Jana G. Delfino 

Division of Imaging, Diagnostics, and Software Reliability, Office of Science and Engineering Laboratories, Center for Devices and Radiological Health, U.S. Food and Drug Administration, Silver Spring, MD 20993, USA

* Correspondence: miguel.lago@fda.hhs.gov

Abstract

Explainability features are intended to provide insight into the internal mechanisms of an Artificial Intelligence (AI) device, but there is a lack of evaluation techniques for assessing the quality of provided explanations. We propose a framework to assess and report explainable AI features in medical images. Our evaluation framework for AI explainability is based on four criteria that relate to the particular needs in AI-enabled medical devices: (1) Consistency quantifies the variability of explanations to similar inputs; (2) plausibility estimates how close the explanation is to the ground truth; (3) fidelity assesses the alignment between the explanation and the model internal mechanisms; and (4) usefulness evaluates the impact on task performance of the explanation. Finally, we developed a scorecard for AI explainability methods in medical imaging that serves as a complete description and evaluation to accompany this type of device. We describe these four criteria and give examples on how they can be evaluated. As a case study, we use Ablation CAM and Eigen CAM to illustrate the evaluation of explanation heatmaps on the detection of breast lesions on synthetic mammographies. The first three criteria are evaluated for task-relevant scenarios. This framework establishes criteria through which the quality of explanations provided by medical devices can be quantified.

Keywords: artificial intelligence; medical imaging; explainability; transparency; interpretability; heatmaps

1. Introduction

Traditionally, machine learning models relied on hand-crafted features designed by developers for specific tasks. These feature-engineered machine learning systems were easy to understand, since relevant features were engraved into the model architecture. This simplicity made such models comprehensible, but their limited capabilities often left much to be desired. In the last few decades, developments of complex neural networks such as Deep Neural Networks (DNN) and generative artificial intelligence (AI) have led to unprecedented levels of accuracy [1,2]. However, the opacity of such models has made it difficult to understand the reasons behind specific predictions or outcomes, creating a “black box” effect in which the inner workings of the model are incomprehensible to human users [3–5].

Artificial intelligence (AI) applications deployed in critical domains, such as healthcare, raise the need for end users to understand, interpret, and comprehend the logic inside “black box” models. A recent workshop held by the American College of Radiology concluded that “clinical implementation of AI in radiology will continue to be limited



Academic Editors: Quincy A. Hathaway and Yashbir Singh

Received: 31 October 2025

Revised: 5 January 2026

Accepted: 10 January 2026

Published: 16 January 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\) license](https://creativecommons.org/licenses/by/4.0/).

until the safety, effectiveness, reliability, and transparency of such tools are more fully addressed" [6]. Further, the American Medical Association emphasizes the need for AI tools that focus on transparency and explainability in their designs (<https://www.ama-assn.org/system/files/ama-ai-principles.pdf> (accessed on 9 January 2026)). However, the lack of a precise definition of explainability complicates the implementation of explainable AI methods. A clear understanding of what explainability entails is important for effective design and implementation.

Explaining a CNN-based model's reasoning is becoming a very popular topic in the literature. Haar et al. [7] provides a comprehensive review of recent gradient, perturbation, and approximation-based explainability methods for CNNs. Gilpin et al. [8] conducts a survey of available literature, providing a taxonomy of definitions and focusing on what is being explained by provided model explanations. This prior work focuses on *development* of explainable AI systems; here our focus is on *evaluation* of explanations provided by AI medical devices.

The topic of evaluating explainability has received significantly less interest in the literature, and the available literature highlights the need for a specialized evaluation framework. Gunning et al. [9] conclude that "*User studies are still the gold standard for assessing explanations*", which further motivates the need for an evaluation framework that reduces the need for such user studies. A recent review on explainable AI techniques in radiology and nuclear medicine from de Vries et al. [10] conclude that "*There is currently no clear consensus on how explainable AI (i.e., XAI or explainability) should be deployed in order to close the gap between medical professionals and DL (Deep Learning) algorithms for clinical implementation.*" Finally, Shifa et al. [11] conclude that "*future research should explore additional evaluation frameworks that consider both the quantitative alignment and the clinical relevance of AI explanations.*"

Useful explanations must balance providing necessary information with avoiding overwhelming users with more information than they can process [12]. Showing too much information without enough guidance can create unnecessary cognitive load, confuse users, and degrade human performance. For example, Fuxin et al. [13] showed that study participants spent an average of 140% more time to process the visualization and answer the questions when shown SAG (Structured Attention Graphs) explanations, compared with traditional heatmaps which did not change user processing time compared to no visual explanations. Modifying their visualization of SAG to include interactive arrows that highlight images that are similar (but not identical) to the query decreased participant response times by 30% while increasing accuracy by 30%. Similarly, explanations providing incorrect information can detract from user performance. In a recent deception study with computer aided marks, reader performance suffered (decreased) when readers were presented with intentionally incorrect computer aided marks [14], underscoring the need to ensure that all information provided to users by AI systems is accurate and beneficial. Other studies have also demonstrated that biased explanations from AI models can lead to decreased diagnostic accuracy [15,16]. These prior works underscore the need for development of a comprehensive framework for evaluating performance of explainability features such as heatmaps, counterfactuals, or confidence scores [17].

For model developers, a comprehensive evaluation framework can offer a means to analyze strengths and weaknesses of explainability methods, enabling model optimization and bench-marking against industry standards. For model evaluators and regulators, a standardized framework allows for objective comparisons, making assessments more efficient and enhancing credibility by aligning with safety and compliance standards. For end users, such as decision makers and doctors in clinical practice, a comprehensive evaluation framework can help foster trust and confidence in AI-outputs and provide insights to

inform clinical choices and support effective patient communication. Providing an explanation of how models arrive at specific recommendations allows doctors to compare the reasoning underlying the AI output against their clinical judgment. As different users have different needs for explainability features, our approach to defining explainability and its evaluation is specifically framed for medical-imaging-relevant tasks from the perspective of the model's end-users (e.g., clinicians, radiologists, physicians, etc.) rather than model developers.

There have been some efforts on creating evaluation guidelines for XAI in medical imaging. For example, Jin et al. [18] propose evaluation guidelines based on understandability, clinical relevance, truthfulness, plausibility, and computational efficiency. However, their evaluation framework involves a human in most metrics, which increases the burden of XAI evaluation for medical device manufacturers and regulators. In a similar sense, Quantus [19] proposes an expandable toolkit to evaluate XAI methods which are based on quantitative metrics for faithfulness, robustness, localization, complexity, randomization, and axiomatic.

We take inspiration from both papers to propose a simplification that relies on the minimum information needed to evaluate XAI, focusing on the specific requirements in medical imaging with the minimal human evaluation. As medical device manufacturers increasingly seek to incorporate explainability features, it is critical to develop a robust, well-established evaluation paradigm.

2. Explaining Explainability

We define *Explainability* as the ability of an AI model to translate its internal reasoning mechanisms in a humanly understandable way. Put another way, explainability refers to a representation of the mechanisms underlying the operation of an AI system [20]. Explainable AI can be achieved in two main ways: embedded directly into the architecture of the AI model (intrinsic explainability or model interpretability) or as a post-hoc method.

An explanation of the model's internal mechanisms should be tailored to a specific task (clinical relevancy) and for a specific expertise. The explanation seeks to describe the process, or rules, that were implemented to achieve a specific model output, independent of the context in which that output will be used [3]. Typically, explanations are detailed, technical, and may be causative. Decisions will be based on a combination of the explanation and the model output. Explainable systems can be debugged and monitored more easily and lend themselves to more thorough documentation, audit, and governance [20].

Explainability in medical imaging can take many forms [21], including: heatmaps (i.e., saliency maps) that highlight regions of the image that contribute to the model response, prototypes that shows examples of similar inputs, counterfactuals that shows a different or opposite class as an explanation, and confidence ratings that shows a score for the likelihood of an output to be correct. This manuscript will focus on evaluating explainability features of Convolutional-Neural-Network-based models.

3. Evaluating Explainability

In this section, we present a framework for evaluating explainability methods and an explainability card for reporting results (see Supplementary Material for an example). This card is divided in two sections: (1) a descriptive section containing information about the explainability method; and (2) a quantitative section that assesses the performance of the method using relevant criteria and metrics.

Such a card can offer clarity and understanding of how different explainability methods perform, enabling informed decision-making, benchmarking, and targeted improvements. It can also ensure quality assurance and facilitate iterative improvement. Addition-

ally, using this card can facilitate improved communication by providing a standardized format for presenting and discussing explainability metrics and methods, ensuring that everyone has access to the same information and can engage in more productive dialogues.

3.1. Descriptive Information

The descriptive information covers everything related to the explainability method and the application including an (1) overview of the method, (2) context of use, (3) limitations and recommendations, and (4) validation setting.

The overview provides the name of the method used to generate the explanation, its description and abbreviation, its type (e.g., local, post hoc), citation or reference, and the software (and version) used to implement the method.

The context of use section identifies the audience, specific task and the model to which the explainability method is applied. Understanding the context of use is important as the effectiveness and relevance of an explainability method can vary depending on the task. It is hard to fully understand and interpret the meaning of the explainability method without knowing this context. For example, a heatmap is typically used to provide visual explanations for classification models by highlighting important regions in an image that contribute to the model output. However, its application would differ in a regression task, where the focus might be on understanding how different regions of the input contribute to a continuous output variable, and in a segmentation task, where the heatmap would need to explain pixel-level decisions rather than overall image classification.

The limitations and recommendations section provides an overview of the constraints, limitations, and assumptions associated with the explainability method to ensure that stakeholders are well-informed about its practical aspects. There may be cases where the method produces misleading results, particularly in edge cases or under atypical conditions. Therefore, understanding these limitations is necessary for correctly interpreting the outputs of the explainability method. Additionally, it is important to recognize if the explainability method relies on certain assumptions, such as the presence of a stable underlying model, and specific operating conditions. This section also lists specific cases or conditions where the method fails or underperforms. By highlighting these failure cases, users can better anticipate potential issues and take preventive measures. In addition to the limitation(s), recommendations can be included to maximize the utility of the explainability method and integrate the method into existing workflows.

Finally, the validation setting section should provide details about the design of the validation process. For example, this section can provide information related to user studies.

3.2. Quantitative Evaluation

To evaluate the performance of explainability methods, we first need to establish a set of evaluation criteria (Figure 1). In this section, we define four key evaluation criteria and present different approaches that can be used to assess their performance. Each criterion is explained in terms of relevance to clinical applications, and we provide demonstrative examples from two explainability methods: Eigen CAM and Ablation CAM.

Eigen CAM [22] and Ablation CAM [23] are two widely used explainability methods for visualizing how deep learning models make decisions. Eigen CAM highlights important regions by computing the principal components of feature maps, while Ablation CAM systematically removes parts of the model to assess the contribution of each component to the final prediction. These methods are complementary, providing both global and local insights into model explainability.

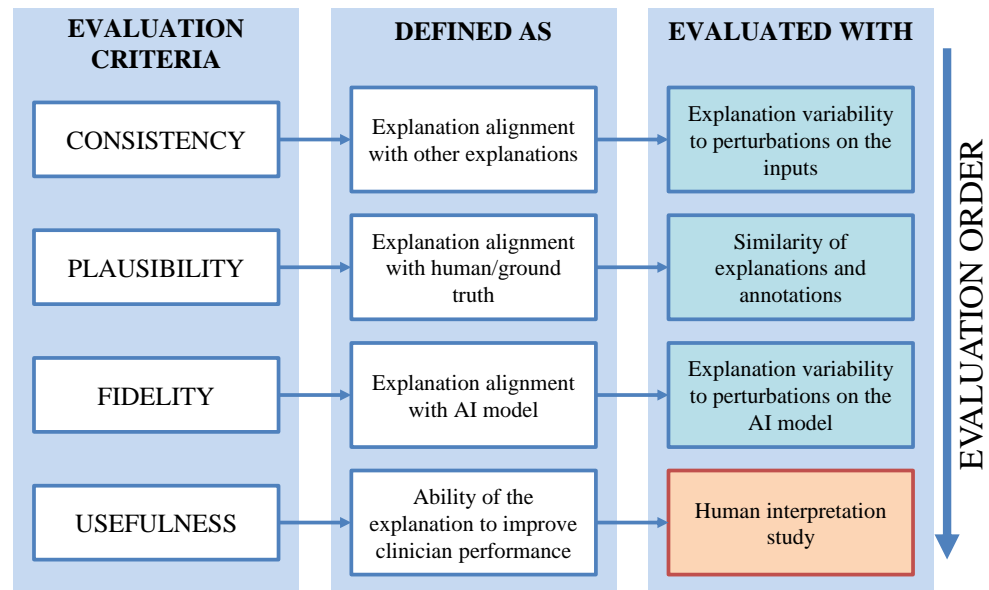


Figure 1. Explainability criteria, their definitions, and their evaluation.

We leveraged the M-SYNTH dataset [24] which includes synthetic x-ray mammography images. The use of synthetic data allows us to have a perfectly known ground-truth (i.e., lesion presence and location) and reduce the variability in clinical images. This allowed us to isolate the analysis of the heatmaps without needing to consider annotations quality in a clinical dataset. More specifically, we used a sample size of 150 images where 50% contain 1 lesion.

We applied Eigen CAM and Ablation CAM to explain the output of a Faster R-CNN model [25] trained to detect lesions on digital mammography datasets. The Faster R-CNN was trained on fatty breast density images with a lesion radius of 7 mm. After training, the model had an area under the receiver-operating characteristic (ROC) curve (AUC) value of 1 for the fatty breast density, with decreasing AUC as density increased (the model had the worst performance with the dense density in the M-SYNTH dataset). We used a smaller lesion (5 mm radius) on our testing dataset to reduce the AUC and add variability to the model performance.

3.2.1. Consistency

Consistency measures how reliably an explainability method provides similar explanations when small variations are introduced to the input data. These variations could be noise, slight image rotations, or shifts, which are often seen in real-world applications such as patient movement during imaging. Consistency is also referred to as stability or robustness.

A consistent visual-based explainability method should generate similar explanations, even when the input image undergoes small changes. Evaluating the consistency of the explanation provided is important as it builds trust and reliability by providing predictable and dependable explanations. Consistency should be the first step toward a good explanation.

To measure consistency, various metrics can be applied to compare explanations provided for the original and transformed images. For example, distance-based metrics can assess the difference between pixel intensities in the heatmaps of the perturbed and non-perturbed images, or compare two feature vectors. Overlap metrics, such as Intersection over Union (IoU), can quantify how much the regions of interest highlighted in the heatmaps overlap between the original and transformed images or image similarity metrics such as SSIM (Structural Similarity Index Measure) or MSE (Mean Squared Error) can be

applied. Consistency can be evaluated by comparing explanations for similar cases under different perturbations, such as noise, rotations, or image degradation. These perturbations allow for the quantification of consistency based on how much the explanation changes under varying degrees and types of alterations. A highly consistent explanation would demonstrate stable metric values across different scenarios and would remain relatively unaffected by such perturbations.

Figure 2 illustrates how the consistency of heatmaps generated using Eigen CAM and Ablation CAM vary under different conditions in digital mammography. Specifically, the figure explores changes in heatmaps due to variations in x-ray radiation doses and degrees of image rotation. In Figure 2A, consistency is quantified using SSIM, MSE, and IoU, with the corresponding values compared across different radiation doses and rotation angles. Despite the model's accuracy remaining stable, SSIM, MSE, and IoU values exhibit notable fluctuations, which highlights how explainability methods can behave inconsistently under subtle changes. Figure 2B presents visual examples of heatmaps, demonstrating how they change when the radiation dose is reduced and when the image is rotated by 50 degrees. The top left shows the heatmap under a high radiation dose, while the bottom left shows the heatmap at a lower dose. Similarly, the top right displays the heatmap for the original image, and the bottom right shows the heatmap after a 50-degree rotation, where both the input and heatmap were set back to their original position for demonstration purposes.

This analysis underscores the need to evaluate heatmap consistency under clinically relevant variations, such as changes in radiation dose or image rotation. If small and insignificant variations in the input data can cause large shifts in the explanation, the method becomes difficult to rely on, as it may produce erratic interpretations under real-world conditions. Therefore, consistency should be considered the first step in evaluating any explainability method.

3.2.2. Plausibility

Plausibility measures the alignment between the generated explanation and ground truth labels provided by human explanations. For example, in visual-based explanations such as heatmaps, a plausible explanation for detecting pneumonia in a medical image should highlight areas in the lungs where a radiologist would expect to see signs of pneumonia. If the heatmap aligns with areas a radiologist identifies as relevant, it is considered plausible. In non-visual-based explanations, such as feature importance scores in a model predicting patient outcomes based on clinical features, a plausible explanation might highlight high blood pressure and cholesterol as important factors for heart disease, aligning with established medical knowledge or reference standard.

To quantitatively assess the alignment between a model's explanation and domain knowledge (ground truth) in a medical image, various overlap metrics can be employed. For example, metrics like IoU can be used to compare heatmaps or localization outputs with ground-truth bounding boxes, segmentation masks, or human attention maps in imaging data. Additionally, correlation metrics such as Spearman's rank correlation and rank correlation [26] can be used to evaluate the correspondence between the generated and ground truth explanations.

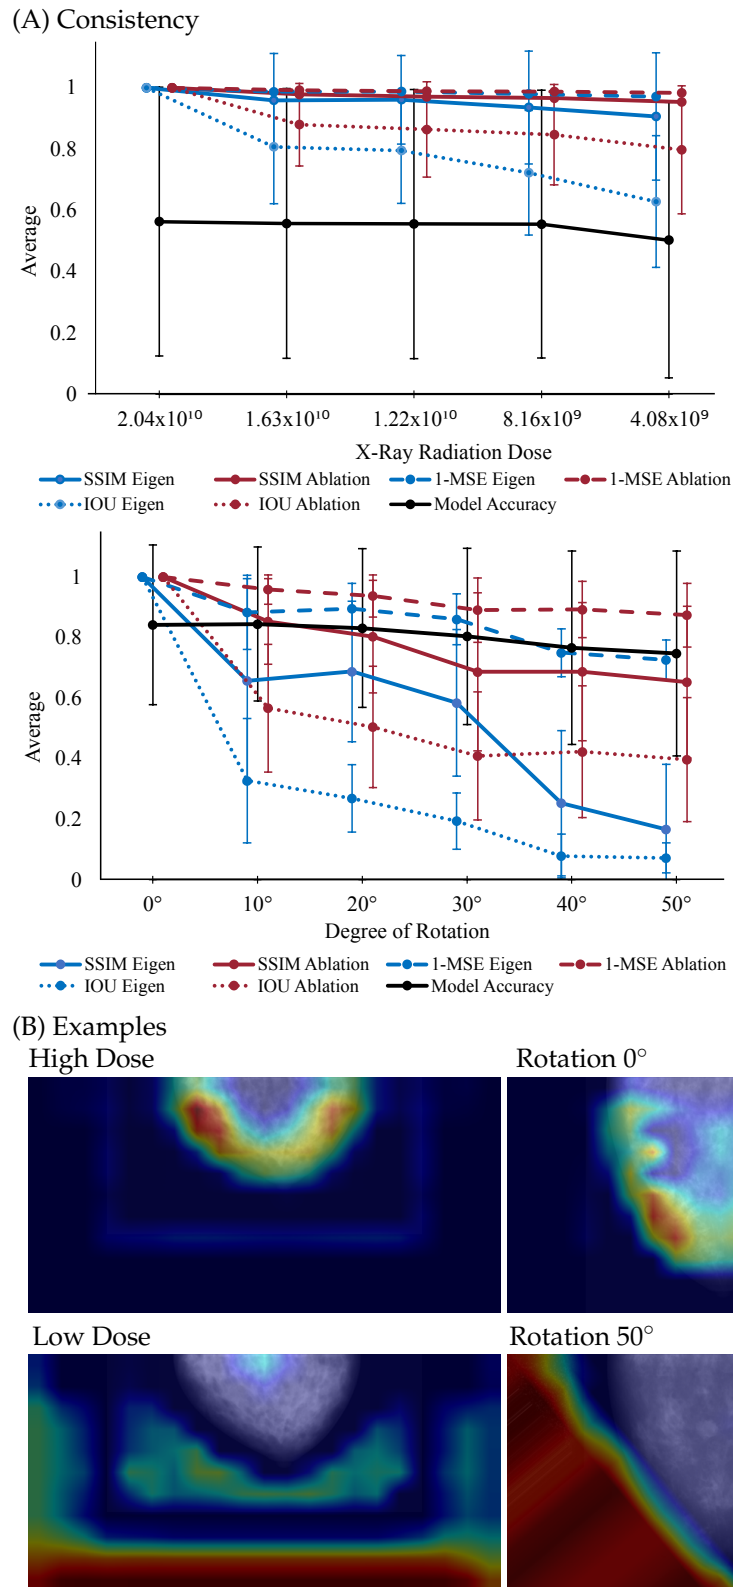


Figure 2. (A) Results for Consistency based on Radiation Dose and Rotation. Metrics are calculated using Structural Similarity Index Measure (SSIM), Mean Squared Error (MSE), and Intersection over Union (IoU). (B) Example of a heatmap from two different doses (left) and two different rotations (right) on a synthetic breast mammography.

Figure 3 illustrates examples of heatmaps with varying levels of plausibility. Figure 3A presents the quantitative plausibility results, where IoU and Model Accuracy were used to calculate the average plausibility scores. As shown in the plot, Ablation CAM demonstrates a signifi-

cantly higher plausibility score compared to Eigen CAM, indicating better alignment with ground truth lesion locations. While this approach captures broader activation patterns, it may not always localize the specific regions most relevant to the clinical task, such as detecting small lesions. In contrast, Ablation CAM systematically removes parts of the model to assess each component's contribution to the prediction, allowing it to focus more on localized regions that directly influence the model's decision.

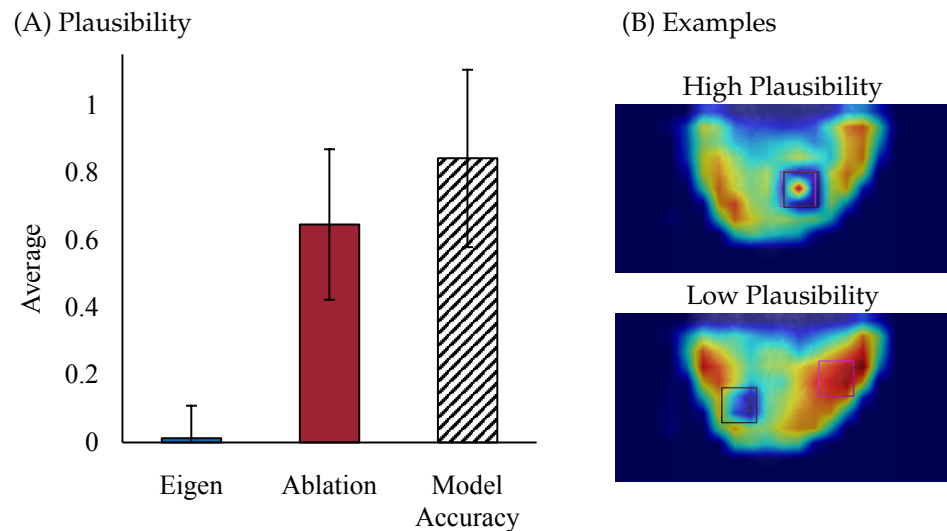


Figure 3. (A) Results for Plausibility. Metrics are calculated using Intersection Over Union (IOU) and Model Accuracy. (B) Examples of high and low plausibility heatmaps. Pink squares are the ROI centered in the highest activation area. Black squares are the ground truth lesion location ROI.

Figure 3B provides visual examples of high and low plausibility heatmaps. In the high plausibility example, the heatmap aligns closely with the ground truth lesion location (represented by the black square), and the region of interest (ROI) in the heatmap (indicated by the pink square) is centered on the lesion, demonstrating good alignment between the model's explanation and human expectations. Conversely, the low plausibility example shows poor alignment between the heatmap's highest activation area (pink square) and the true lesion location (black square), with the highlighted region failing to correspond accurately to the lesion. These results show that some explainability methods can be less plausible, even if they are consistent. Therefore, after assessing the consistency of explainability methods, it is important to ensure they provide explanations that align with human ground truth.

A comprehensive evaluation of plausibility may require incorporating different levels of ground truth. For example, in mammography, the most relevant ground truth might be the lesion region itself, as annotated by expert radiologists. At the same time, surrounding anatomical structures or contextual imaging features, also constitutes valuable information, as humans often consider these regions when making decisions. This information can be obtained through tools like eye-tracking, which capture whether radiologists focus not only on lesion regions but also on adjacent areas that provide contextual clues about a lesion's characteristics. By incorporating this hierarchy of ground truth—from precise lesion annotations to indirect contextual cues—a more nuanced and comprehensive evaluation of plausibility can be achieved.

3.2.3. Fidelity

Fidelity, also referred to as faithfulness or truthfulness, measures the agreement of an explanation with the underlying mechanisms the model used to predict an outcome. A

high-fidelity explanation accurately reflects the factors the model considered important in its decision. For example, in an image-based task, the pixels representing the detected object should be highlighted for the model to make its prediction, while non-relevant pixels should remain unhighlighted. It is important to note that model's accuracy and fidelity are related but distinct concepts. Fidelity refers to how well the explanation aligns with the internal workings of the model, while accuracy measures how well the model performs at its task. If a model is highly accurate, and the explanation has high fidelity, this means the explanation accurately reflects the true reasons behind the model's predictions. However, some explanations may only provide local fidelity, meaning they accurately represent the model's reasoning for a specific subset of the data (e.g., a particular instance or a small group of examples). For example, local surrogate models or Shapley Values explain individual predictions by focusing on how the model behaves for that particular instance, rather than providing a global understanding of the entire model's decision-making process.

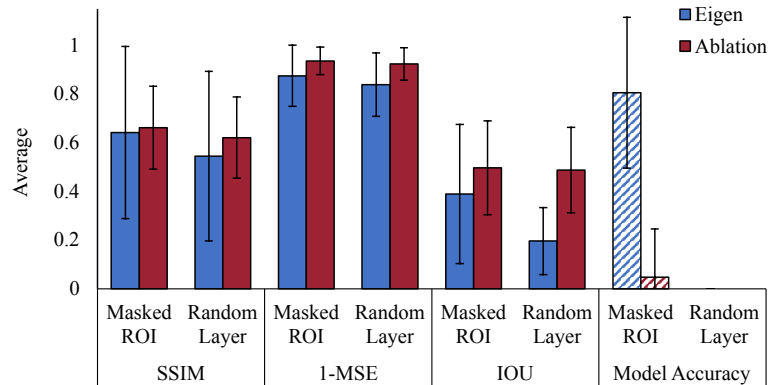
Several approaches exist to measure the fidelity of an explainability method. One well-known approach, introduced by [27], is the Model Parameter Randomization Check. In this method, the model's parameters are perturbed by randomizing or re-initializing the weights, which disrupts the original learned patterns. Explanations are then generated both before and after this randomization process. The key idea is to compare the explanations: explanations that remain the same after randomization suggest that the explanation is not sensitive to the model's learned parameters and does not faithfully reflect the model's reasoning. On the other hand, explanations that differ after randomization indicate that the explanation is tied to the model's internal mechanisms. Another method for assessing fidelity is the White Box Check [28], which evaluates explainability methods by using a transparent model with known decision-making processes, such as a decision tree or random forest. The generated explanations are compared to the model's internal reasoning, allowing for a direct assessment of how accurately the explanation method reflects the model's true decision-making process.

Additional methods include the Single Deletion Method and the Incremental Deletion Method [28,29]. The Single Deletion Method evaluates feature importance by testing how much the model's output changes when a single feature is removed or altered. A correct explanation should show that removing the most important feature leads to the largest change in the model's output, while removing less important features causes minimal impact. The Incremental Deletion Method extends this by removing features one by one, either in order of importance or in reverse, to assess how the model's output changes. This approach helps verify whether the cumulative removal of features corresponds to their ranked importance.

Figure 4 illustrates the fidelity of two saliency map methods. As shown in Figure 4A, Ablation CAM consistently exhibits higher fidelity than Eigen CAM, indicating that it more faithfully reflects the model's internal reasoning when the model's parameters are perturbed. Particularly, with each metric, Ablation CAM has reduced similarity to the un-perturbed explanation, but correspondingly reduces the accuracy of the model. Together, this suggests that Ablation CAM is producing higher fidelity explanations due to disruption in similarity and model performance. In contrast, Eigen CAM shows lower fidelity, with increased variability between the different metrics and a spared model accuracy after perturbation for the masked ROI case, which highlights its weaker connection to the model's internal mechanisms. Together, these results suggest that comparison between explanations before and after perturbation is necessary to examine fidelity but that examining accuracy is also critical for a full account of fidelity. Though fidelity and model accuracy are independent, understanding how the explanation aligns with the model's decision is critical to fully understanding fidelity. The pattern of results presented in Figure 4 is a great

example of this need for both pieces of information, as both explanations were disrupted by perturbations, but only Ablation CAM resulted in a change in the model’s decision making. However, if just examining model accuracy, in the case of randomizing model layers, we would conclude that both explanations have high fidelity, which is not the case.

(A) Fidelity



(B) Examples

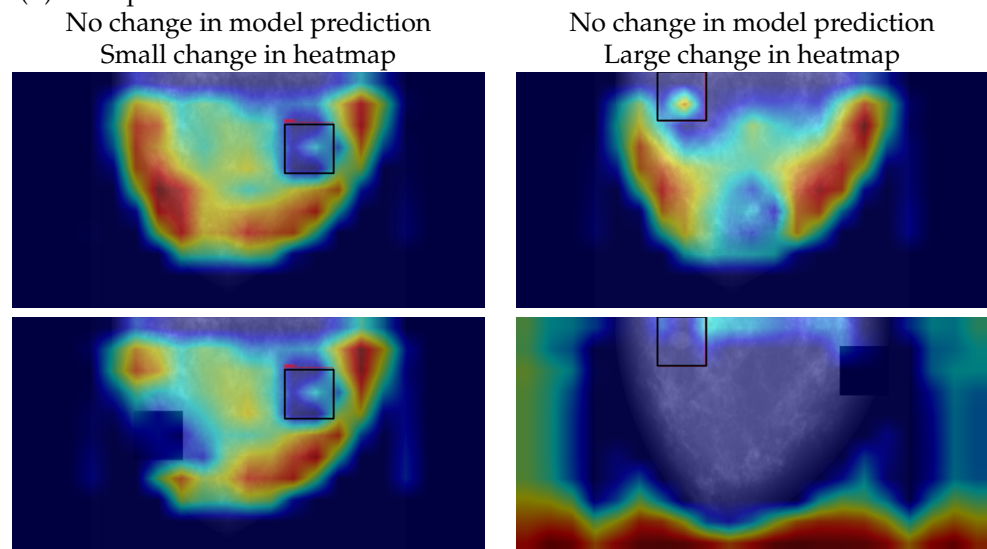


Figure 4. (A) Results for Fidelity. Metrics are calculated using Structural Similarity Image Metric (SSIM), Mean Square Error (MSE), Intersection Over Union (IOU), and Model Accuracy. (B) Examples of low (left) and high (right) fidelity heatmaps.

Figure 4B offers visual examples that further illustrate these findings when we removed the highest activation ROI from the original input (top) and fed it back through the model (bottom). On the left column, the example shows no change in the model prediction and only a small change in the heatmap, indicating that the explanation is not affected even after minor perturbations to the input, which reflects lower fidelity. On the right column, another example shows no change in the model prediction but a large change in the heatmap, suggesting that the explanation is sensitive to perturbations based on the highest heatmap activation areas, indicating higher fidelity.

It is important to note the difference between plausibility and fidelity. Plausibility focuses on measuring the alignment between the explanation and domain knowledge or ground truth, while fidelity focuses on measuring the alignment of the explanation with respect to the predictive model. For example, an explanation highlighting patient information in the background to distinguish between sick and healthy patients would score low on plausibility with respect to tumor segmentation masks but could still correctly

reflect the reasoning of this flawed classifier (high fidelity). Therefore, it is important to evaluate explainability methods for both plausibility and fidelity independently.

3.2.4. Usefulness

The final evaluation component focuses on assessing an explanation's ability to help the user in effectively performing a given task. This involves two aspects: Clinical Relevance and Interpretability.

Clinical relevance refers to the degree to which an explanation is useful and informative for the specific task it is intended to support. It goes beyond whether the explanation merely aligns with general ground truth (as in plausibility) and instead focuses on whether the explanation directly addresses the critical aspects of the clinical task. This criterion evaluates if the information provided by the model is meaningful and necessary for clinicians to make informed decisions. For example, in a breast density evaluation task, clinical relevance would require that the model's explanation focuses on the dense tissue areas in a mammogram, as these regions are essential for assessing breast density. An explanation that highlights irrelevant or peripheral areas may still seem plausible in some contexts, but would lack clinical relevance because it does not inform the task at hand. Thus, clinical relevance ensures that the explanation contributes directly to the decision-making process and enhances the clinician's ability to interpret the model's output in a medically useful way.

Interpretability refers to the extent to which a non-AI expert (e.g., a clinician such as a radiologist) can extract and understand information from the provided explanation. It is often synonymous with understandability and concerns how easily the internal workings of the model can be made accessible and understandable to end users. The National Institute of Standards and Technology (NIST) attempts to draw a distinction between interpretation (the process of deriving meaning from a stimulus) and explanation (the process of generating detailed description of how the outcome was reached [3]), and we generally agree with this distinction. Interpretability focuses on how well the model is designed to be understood by human users. A model can still be difficult to interpret if the information processing is too complex for the user to comprehend (e.g., deep neural network).

The concepts of clinical relevance and interpretability are related, as both focus on how well an AI-generated explanation can assist human users, especially non-AI experts, in understanding and making decisions based on the model's output. Clear and easy-to-understand explanations (high interpretability) are more likely to be clinically relevant, as they reduce cognitive load and enhance decision-making. However, an explanation can be interpretable but not clinically relevant if it highlights the wrong areas or fails to assist with the specific task. Similarly, a clinically relevant explanation can be hard to interpret if the model's inner workings are too complex. Low interpretability can negatively impact a user's performance, thus reducing trust in the system. Superficial or unrelated explanations can increase a user's cognitive load and decrease the efficiency and/or performance of the user.

Assessing usefulness is only possible through studies in which medical professionals evaluate the generated explanations. The human evaluation should focus on interpretability [30–33] or time-to-decision improvement [34]. One possibility is that users are asked to perform the selected task both with and without the provided explanation; if the provided explanation is useful, it should either enable enhanced performance on the task or completion of the task more quickly.

3.3. From Quantitative to Human-Based Assessment

We have defined four criteria for evaluating explainability methods: Consistency, Plausibility, Fidelity, and Usefulness. The order of evaluation is important. Consistency

must be evaluated first, as it ensures that the explainability method provides stable and reproducible explanations across similar inputs and conditions. Once consistency is confirmed, plausibility is assessed to measure how well the explanation aligns with expert annotations. By evaluating plausibility early, we ensure that the explanation aligns with human perspective. If the explanation fails this test, it may lack meaningful insight, regardless of its connection to the model's internal mechanisms. After plausibility, fidelity is assessed to determine whether the explanation accurately reflects the model's reasoning. Even if an explanation is plausible to human experts, it can still be misleading if it does not correctly represent how the model arrived at its decision.

Once an explainability method passes Consistency, Plausibility, and Fidelity, it is ready to move into the second stage for human-based evaluation, where medical experts assess the explanation's usefulness for specific clinical tasks. Requiring a method to pass consistency, plausibility, and fidelity before involving human experts, filters out poor or misleading explanations, ensuring that only a subset of explanations require human evaluation.

4. Discussing Explainability

While explainability features are desired by users of AI-enabled devices, a consensus definition of explainability as well as evaluation techniques for assessing the quality of provided explanations in medical imaging is lacking. This lack of consensus definitions and techniques for assessing explanations is problematic, as a growing body of evidence suggests that providing untrustworthy information to users can degrade task performance [13–16,35,36]. In this work, we propose a framework for the evaluation of explainable AI features. We demonstrate the feasibility of our proposed framework by applying it to heatmap-based explanations in synthetic breast mammography. However, our intention is for this framework to be broadly applicable to different explainable AI methods across a wide range of medical imaging.

Three of our proposed criteria—Consistency, Plausibility, and Fidelity—can be evaluated on previously-collected data, while Usefulness require a human task-based performance assessment. To best use resources, we have structured the evaluation framework in a specific order, proposing the evaluation of Consistency, Plausibility, and Fidelity in that order, followed by Usefulness. This ordering helps to minimize resources, as explanations with low Consistency can be evaluated before moving to Plausibility and Fidelity, which can then proceed to human task-based evaluation studies for Usefulness.

To the best of our knowledge, no clear consensus exists on what makes an explanation 'good' or 'useful.' The ambiguity of this question is at least partially explained by the fact that the interpretation of an explanation is linked to the context in which that explanation is provided. In the current work, we have framed a "good" explanation as one that consistently provides the same information, is aligned with the clinical ground truth, changes with the output of the model, and is useful for the given task. This framing is consistent with prior thinking that an ideal explanation must convey the *gist* of what is needed, but without overwhelming with irrelevant information [3]. Our work extends prior work that explanations should display fidelity, interpretability, and be independent of any specific AI model [37].

Evaluating explanations is confounded by the fact that different users will have different needs, abilities and desires from explanations [37,38]. Put another way, explanations are useful to different groups for different reasons, and one can further argue that explanations are most useful when they are tailored to the needs of the intended user [39]. AI developers generally seek explanations to debug or improve their models. For the developer, a good explanation is causal and helps illustrate why something is occurring (e.g., the algorithm shows bias towards a particular group because the training data was imbalanced), so that

additional action can be taken (e.g., collecting more training data). A user interacting with the same model would be more interested in whether or not the model output could be relied upon as fair, and their potential actions would be to use or not use the model output. Additionally, what can be a good explanation for one task such as breast cancer detection does not necessarily hold true for real-time polyp detection in endoscopy [40] or neurological disorder diagnostics [41]. In this work we have focused on the value of the explanation for end-users for clinically-relevant tasks.

AI-provided explanations must be interpretable by human users. A significant hurdle to achieving this is that humans and computers process information differently. Whereas humans rely on simple, imprecise inclinations to make decisions, machine learning models rely on programmatic verbatim processes to generate predictions [3,42]. Translating the precise technical decisions of a computer into human-interpretable explanations is an ongoing challenge. Additionally, the vast majority of AI-models deployed in healthcare today are used as assistive devices, meaning that the impact of the model is reliant on the combination of the AI model output and the human interpretation. This adds an additional layer of complexity, as users differ in the degree to which they are willing and able to utilize their own background knowledge to interpret detailed technical information, and users place different amounts of trust in AI output [43]. In this area, we can find studies about the effect of explanations on clinical performance [3,31,44,45]. This is an active research area that will be considered for future projects.

5. Limitations

In this paper, we use synthetic mammography data to illustrate the utility of the proposed framework. The use of synthetic data allows us to reduce the number of unknown variables as well as to control the relevant aspects of the task. For instance, having a near-perfect ground truth provides an ideal evaluation of plausibility, as opposed to having the challenge of imperfect or high-variability reference standards [46]. Furthermore, controlling radiological dose permits for studies that would otherwise be unethical and impractical (e.g., testing for consistency with real patients at different dose levels). We intend to explore application of the framework to clinical data in future work.

The metrics used here are only one example of possible ways to quantitatively evaluate explainability features. There are multiple metrics in addition to SSIM, MSE and IOU that can be explored. Our framework aims to be flexible enough to accept other metrics applied to the same criterion. Interpretation of such metrics must be well established for the specific AI and task at hand. The proposed scorecard in the Supplementary Material should explain the meaning and weight of each one of the criteria to avoid incorrect claims.

We acknowledge that some criteria are based on perturbations that also affect the model accuracy. This accuracy loss can be due to several factors and cannot be attributed only to a change in reasoning (for fidelity). This is an assumption on our fidelity metric that will be explored in the future. However, in this work the use of synthetic images allow us to create the same image with and without a lesion, allowing us to isolate the effect of the lesion presence on the provided explanation.

It is noteworthy to mention that this framework does not provide a threshold or any acceptability criteria; instead, the framework offers enough flexibility to include acceptance thresholds for each of the four criteria. These thresholds must be defined and justified independently for different applications, tasks, and imaging modalities. Some AI algorithms, depending on the task, may prioritize high consistency whereas others may focus on better fidelity.

Finally, heatmaps are not the only type of explanation available in medical imaging. Other XAI methods may include counter-factuals, attribution values (SHAP), or concept-

based attributions. Although the framework we propose here is centered around heatmaps, it is generalizable enough that the four main criteria can serve as the starting point for new directed frameworks.

6. Conclusions

Explanations provided by AI systems are seen as a means of increasing trust and confidence in the AI systems. However, to be useful, such explanations must provide only useful information that is interpretable to the user. To fill this gap, we propose a framework for the evaluation of explainability features in AI-models applied to medical data. Our framework is based on a hierarchy of four criteria that evaluate, in order: (1) consistency of explanations, (2) plausibility of explanations, (3) fidelity of explanations, as well as (4) usefulness.

Disclaimer

This article reflects the views of the authors and does not represent the views or policy of the U.S. Food and Drug Administration, the Department of Health and Human Services, or the U.S. Government. The mention of commercial products, their sources, or their use in connection with material reported herein is not to be construed as either an actual or implied endorsement of such products by the Department of Health and Human Services. This is a contribution of the U.S. Food and Drug Administration and is not subject to copyright.

Supplementary Materials: The following supporting information can be downloaded at: <https://www.mdpi.com/article/10.3390/bioengineering13010111/s1>, File S1: XAI Scorecard.

Author Contributions: Conceptualization, M.A.L., G.Z. and B.E.; Methodology, M.A.L., G.Z. and B.E.; Data Curation, B.E.; drafting M.A.L., G.Z., B.E. and J.G.D.; supervision, J.G.D. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: M-SYNTH data is available in <https://huggingface.co/datasets/didsr/msynth> (accessed on 9 January 2026). The data that support the findings of this study are available from the corresponding author upon reasonable request.

Acknowledgments: Brandon Eich acknowledges funding by appointment to the Research Participation Program at the Center for Devices and Radiological Health administered by the Oak Ridge Institute for Science and Education through an interagency agreement between the US Department of Energy and the US Food and Drug Administration.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Guidotti, R.; Monreale, A.; Ruggieri, S.; Turini, F.; Giannotti, F.; Pedreschi, D. A survey of methods for explaining black box models. *ACM Comput. Surv. (CSUR)* **2018**, *51*, 1–42. [[CrossRef](#)]
2. He, C.; Ma, M.; Wang, P. Extract interpretability-accuracy balanced rules from artificial neural networks: A review. *Neurocomputing* **2020**, *387*, 346–358. [[CrossRef](#)]
3. Broniatowski, D.A. *Psychological Foundations of Explainability and Interpretability in Artificial Intelligence*; NIST Technical Report; US Department of Commerce, National Institute of Standards and Technology: Gaithersburg, MD, USA, 2021.
4. Doshi-Velez, F.; Kim, B. Towards A Rigorous Science of Interpretable Machine Learning. *arXiv* **2017**, arXiv:1702.08608. [[CrossRef](#)]
5. Schmidt, P.; Biessmann, F. Quantifying Interpretability and Trust in Machine Learning Systems. *arXiv* **2019**, arXiv:1901.08558. [[CrossRef](#)]

6. Larson, D.B.; Tomkins, K.G.; Zacharias-Andrews, K.; Irani, N.; Pittman, S.M.; Purysko, A.S.; Wandtke, B.; Bhargavan-Chatfield, M. The ACR learning network: Facilitating local performance improvement through shared learning. *J. Am. Coll. Radiol.* **2023**, *20*, 369–376. [[CrossRef](#)]
7. Haar, L.V.; Elvira, T.; Ochoa, O. An analysis of explainability methods for convolutional neural networks. *Eng. Appl. Artif. Intell.* **2023**, *117*, 105606. [[CrossRef](#)]
8. Gilpin, L.H.; Bau, D.; Yuan, B.Z.; Bajwa, A.; Specter, M.; Kagal, L. Explaining Explanations: An Overview of Interpretability of Machine Learning. In Proceedings of the 2018 IEEE 5th International Conference on Data Science and Advanced Analytics (DSAA), Turin, Italy, 1–3 October 2018; pp. 80–89. [[CrossRef](#)]
9. Gunning, D.; Vorm, E.; Wang, J.Y.; Turek, M. DARPA’s explainable AI (XAI) program: A retrospective. *Appl. AI Lett.* **2021**, *2*, e61. [[CrossRef](#)]
10. de Vries, B.M.; Zwezerijnen, G.J.; Burchell, G.L.; van Velden, F.H.; Menke-van der Houven van Oordt, C.W.; Boellaard, R. Explainable artificial intelligence (XAI) in radiology and nuclear medicine: A literature review. *Front. Med.* **2023**, *10*, 1180773. [[CrossRef](#)]
11. Shifa, N.; Saleh, M.; Akbari, Y.; Al Maadeed, S. A review of explainable AI techniques and their evaluation in mammography for breast cancer screening. *Clin. Imaging* **2025**, *123*, 110492. [[CrossRef](#)]
12. Venkatesh, K.; Mutasa, S.; Moore, F.; Sulam, J.; Yi, P.H. Gradient-Based Saliency Maps Are Not Trustworthy Visual Explanations of Automated AI Musculoskeletal Diagnoses. *J. Imaging Inform. Med.* **2024**, *37*, 2490–2499. [[CrossRef](#)]
13. Fuxin, L.; Qi, Z.; Khorram, S.; Shitole, V.; Tadepalli, P.; Kahng, M.; Fern, A. From heatmaps to structured explanations of image classifiers. *Appl. AI Lett.* **2021**, *2*, e46. [[CrossRef](#)]
14. Bernstein, M.H.; Atalay, M.K.; Dibble, E.H.; Maxwell, A.W.; Karam, A.R.; Agarwal, S.; Ward, R.C.; Healey, T.T.; Baird, G.L. Can incorrect artificial intelligence (AI) results impact radiologists, and if so, what can we do about it? A multi-reader pilot study of lung cancer detection with chest radiography. *Eur. Radiol.* **2023**, *33*, 8263–8269. [[CrossRef](#)] [[PubMed](#)]
15. Jabbour, S.; Fouhey, D.; Shepard, S.; Valley, T.S.; Kazerooni, E.A.; Banovic, N.; Wiens, J.; Sjoding, M.W. Measuring the impact of AI in the diagnosis of hospitalized patients: a randomized clinical vignette survey study. *JAMA* **2023**, *330*, 2275–2284. [[CrossRef](#)]
16. Khera, R.; Simon, M.A.; Ross, J.S. Automation Bias and Assistive AI: Risk of Harm From AI-Driven Clinical Decision Support. *JAMA* **2023**, *330*, 2255–2257. [[CrossRef](#)]
17. Klein, L.; Lüth, C.; Schlegel, U.; Bungert, T.; El-Assady, M.; Jäger, P. Navigating the maze of explainable ai: A systematic approach to evaluating methods and metrics. *Adv. Neural Inf. Process. Syst.* **2024**, *37*, 67106–67146.
18. Jin, W.; Li, X.; Fatehi, M.; Hamarneh, G. Guidelines and evaluation of clinical explainable AI in medical image analysis. *Med. Image Anal.* **2023**, *84*, 102684. [[CrossRef](#)]
19. Hedström, A.; Weber, L.; Krakowczyk, D.; Bareeva, D.; Motzkus, F.; Samek, W.; Lapuschkin, S.; Höhne, M.M.M. Quantus: An Explainable AI Toolkit for Responsible Evaluation of Neural Network Explanations and Beyond. *J. Mach. Learn. Res.* **2023**, *24*, 1–11.
20. NIST AI. *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*; The United States Department of Commerce: Washington, DC, USA, 2023.
21. Hossain, M.I.; Zamzmi, G.; Mouton, P.R.; Salekin, M.S.; Sun, Y.; Goldgof, D. Explainable AI for Medical Data: Current Methods, Limitations, and Future Directions. *ACM Comput. Surv.* **2023**, *57*, 1–46. [[CrossRef](#)]
22. Muhammad, M.B.; Yeasin, M. Eigen-CAM: Class Activation Map using Principal Components. In Proceedings of the 2020 International Joint Conference on Neural Networks (IJCNN), Glasgow, UK, 19–24 July 2020. [[CrossRef](#)]
23. Desai, S.; Ramaswamy, H.G. Ablation-CAM: Visual Explanations for Deep Convolutional Network via Gradient-free Localization. In Proceedings of the 2020 IEEE Winter Conference on Applications of Computer Vision (WACV), Snowmass Village, CO, USA, 1–5 March 2020; pp. 972–980. [[CrossRef](#)]
24. Sizikova, E.; Saharkhiz, N.; Sharma, D.; Lago, M.; Sahiner, B.; Delfino, J.; Badano, A. Knowledge-based in silico models and dataset for the comparative evaluation of mammography AI for a range of breast characteristics, lesion conspicuities and doses. *Adv. Neural Inf. Process. Syst.* **2024**, *36*, 37401–37412.
25. Ren, S.; He, K.; Girshick, R.; Sun, J. Faster R-CNN: Towards real-time object detection with region proposal networks. *Adv. Neural Inf. Process. Syst.* **2015**, *28*, 91–99. [[CrossRef](#)] [[PubMed](#)]
26. Vilone, G.; Longo, L. Notions of explainability and evaluation approaches for explainable artificial intelligence. *Inf. Fusion* **2021**, *76*, 89–106. [[CrossRef](#)]
27. Adebayo, J.; Gilmer, J.; Muelly, M.; Goodfellow, I.; Hardt, M.; Kim, B. Sanity checks for saliency maps. *Adv. Neural Inf. Process. Syst.* **2018**, *31*, 9505–9515.
28. Nauta, M.; Trienes, J.; Pathak, S.; Nguyen, E.; Peters, M.; Schmitt, Y.; Schlötterer, J.; Van Keulen, M.; Seifert, C. From anecdotal evidence to quantitative evaluation methods: A systematic review on evaluating explainable AI. *ACM Comput. Surv.* **2023**, *55*, 1–42. [[CrossRef](#)]
29. Li, Y.; Hao, W.; Zeng, H.; Wang, L.; Xu, J.; Routray, S.; Jhaveri, R.H.; Gadekallu, T.R. Cross-Scale Texture Supplementation for Reference-based Medical Image Super-Resolution. *IEEE J. Biomed. Health Inform.* **2025**. [[CrossRef](#)]

30. Ribeiro, M.T.; Singh, S.; Guestrin, C. Anchors: High-precision model-agnostic explanations. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence and Thirtieth Innovative Applications of Artificial Intelligence Conference and Eighth AAAI Symposium on Educational Advances in Artificial Intelligence*; AAAI Press: Washington, DC, USA, 2018; AAAI'18/IAAI'18/EAAI'18.
31. Răz, T.; Pahud De Mortanges, A.; Reyes, M. Explainable AI in medicine: Challenges of integrating XAI into the future clinical routine. *Front. Radiol.* **2025**, *5*, 1627169. [[CrossRef](#)] [[PubMed](#)]
32. McNamara, S.L.; Yi, P.H.; Lotter, W. The clinician-AI interface: Intended use and explainability in FDA-cleared AI devices for medical image interpretation. *NPJ Digit. Med.* **2024**, *7*, 80. [[CrossRef](#)]
33. Kibria, M.G.; Kucirka, L.; Mostafa, J. Assessing AI Explainability: A Usability Study Using a Novel Framework Involving Clinicians. In *Proceedings of the 2025 IEEE 13th International Conference on Healthcare Informatics (ICHI)*, Rende, Italy, 18–21 June 2025; pp. 553–564.
34. Lakkaraju, H.; Bach, S.H.; Leskovec, J. Interpretable decision sets: A joint framework for description and prediction. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, San Francisco, CA, USA, 13–17 August 2016; pp. 1675–1684.
35. Abgrall, G.; Holder, A.L.; Chelly Dagdia, Z.; Zeitouni, K.; Monnet, X. Should AI models be explainable to clinicians? *Crit. Care* **2024**, *28*, 301. [[CrossRef](#)] [[PubMed](#)]
36. Makridis, G.; Fatouros, G.; Kiourtis, A.; Kotios, D.; Koukos, V.; Kyriazis, D.; Soldatos, J. Towards a unified multidimensional explainability metric: Evaluating trustworthiness in AI models. In *Proceedings of the 2023 19th International Conference on Distributed Computing in Smart Systems and the Internet of Things (DCOSS-IoT)*, Pafos, Cyprus, 19–21 June 2023; pp. 504–511.
37. Celino, I. *Who Is This Explanation for? Human Intelligence and Knowledge Graphs for eXplainable AI*; IOS Press: Amsterdam, The Netherlands, 2020. [[CrossRef](#)]
38. Maxwell, W.; Beaudouin, V.; Bloch, I.; Bounie, D.; Cléménçon, S.; d'Alché Buc, F.; Eagan, J.; Mozharovskiy, P.; Parekh, J. Identifying the 'Right' Level of Explanation in a Given Situation. In *Proceedings of the First International Workshop on New Foundations for Human-Centered AI (NeHuAI)*, Santiago de Compostella, Spain, 4 September 2020; p. 63.
39. Combi, C.; Amico, B.; Bellazzi, R.; Holzinger, A.; Moore, J.H.; Zitnik, M.; Holmes, J.H. A manifesto on explainability for artificial intelligence in medicine. *Artif. Intell. Med.* **2022**, *133*, 102423. [[CrossRef](#)] [[PubMed](#)]
40. Huang, C.W.; Su, C.C.; Chou, C.K.; Mukundan, A.; Karmakar, R.; Chen, T.H.; Shukla, P.; Gupta, D.; Wang, H.C. Software-Based Transformation of White Light Endoscopy Images to Hyperspectral Images for Improved Gastrointestinal Disease Detection. *Diagnostics* **2025**, *15*, 1664. [[CrossRef](#)]
41. Hafeez, Y.; Memon, K.; Al-Quraishi, M.S.; Yahya, N.; Elferik, S.; Ali, S.S.A. Explainable AI in diagnostic radiology for neurological disorders: A systematic review, and what doctors think about it. *Diagnostics* **2025**, *15*, 168. [[CrossRef](#)]
42. Arrieta, A.B.; Díaz-Rodríguez, N.; Del Ser, J.; Bennetot, A.; Tabik, S.; Barbado, A.; García, S.; Gil-López, S.; Molina, D.; Benjamins, R.; et al. Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Inf. Fusion* **2020**, *58*, 82–115. [[CrossRef](#)]
43. Ehrmann, D.E.; Gallant, S.N.; Nagaraj, S.; Goodfellow, S.D.; Eytan, D.; Goldenberg, A.; Mazwi, M.L. Evaluating and reducing cognitive load should be a priority for machine learning in healthcare. *Nat. Med.* **2022**, *28*, 1331–1333. [[CrossRef](#)] [[PubMed](#)]
44. Gulum, M.A.; Trombley, C.M.; Ozen, M.; Esen, E.; Aksamoglu, M.; Kantardzic, M. Why are explainable AI methods for prostate lesion detection rated poorly by radiologists? *Appl. Sci.* **2024**, *14*, 4654. [[CrossRef](#)]
45. Rezaeian, O.; Bayrak, A.E.; Asan, O. Explainability and AI confidence in clinical decision support systems: Effects on trust, diagnostic performance, and cognitive load in breast cancer care. *Int. J. Hum. Comput. Interact.* **2025**, 1–21. [[CrossRef](#)]
46. Song, Y.; Liu, Y.; Lin, Z.; Zhou, J.; Li, D.; Zhou, T.; Leung, M.F. Learning from AI-generated annotations for medical image segmentation. *IEEE Trans. Consum. Electron.* **2024**, *71*, 1473–1481. [[CrossRef](#)]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.