



# Explainable AI (XAI) in image segmentation in medicine, industry, and beyond: A survey

Rokas Gipiškis<sup>a,\*</sup>, Chun-Wei Tsai<sup>b</sup>, Olga Kurasova<sup>a</sup>

<sup>a</sup> Vilnius University, Institute of Data Science and Digital Technologies, Lithuania

<sup>b</sup> National Sun Yat-sen University, Department of Computer Science and Engineering, Taiwan

Received 25 June 2024; received in revised form 29 August 2024; accepted 12 September 2024

Available online xxx

## Abstract

Explainable AI (XAI) has found numerous applications in computer vision. While image classification-based explainability techniques have garnered significant attention, their counterparts in semantic segmentation have been relatively neglected. Given the prevalent use of image segmentation, ranging from medical to industrial deployments, these techniques warrant a systematic look. In this paper, we present the first comprehensive survey on XAI in semantic image segmentation. We analyze and categorize the literature based on application categories and domains, as well as the evaluation metrics and datasets used. We also propose a taxonomy for interpretable semantic segmentation, and discuss potential challenges and future research directions.

© 2024 The Author(s). Published by Elsevier B.V. on behalf of The Korean Institute of Communications and Information Sciences. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

**Keywords:** XAI; Interpretable AI; Interpretability; Image segmentation; Semantic segmentation

## 1. Introduction

In the past decade, Artificial Intelligence (AI) systems have achieved impressive results, most notably in natural language processing and computer vision. The performance of such systems is typically measured by evaluation metrics that vary depending on the task but aim to assess the system's outputs. Today's leading AI systems largely rely on deep learning (DL) models, multi-layered neural networks that tend to exhibit increasingly complicated structures in terms of model parameters. The growing complexity of such systems resulted in them being labeled as "black boxes". This highlights that the evaluation metric does not show the full picture: even if its measurement is correct, it does not give insights into the inner workings of the model.

The field of explainable AI (XAI) encompasses different branches of methods that attempt to give insights into a model's inner workings, explain outputs, or make the entire system more interpretable to end users, such as human decision-makers. There is ongoing debate regarding XAI terminology. Concepts like interpretability, explainability, understanding, reasoning, and trustworthiness are challenging

to formalize. While some authors use "interpretable" and "explainable" interchangeably [1], others distinguish between the two [2,3]. When the distinction is made, it is usually to demarcate post-hoc explanations, a type of XAI techniques applied to the already-trained model, and inherently interpretable models [2]. This way interpretability becomes associated with the transparency of the model itself and depends on the ease with which one can interpret the model. For instance, a simple decision tree-based model might be considered more interpretable than a DL model composed of millions of parameters, provided that the former is not too deep. Explainability, in contrast, is often limited to understanding the model's results rather than the model as a whole. While we acknowledge such distinction, throughout this survey we will use "interpretable" and "explainable" synonymously, reserving more specific "architecture-based" and "inherently interpretable" terms when discussing model-specific XAI modifications. This is because not many of the surveyed papers use the term interpretability in a second sense. Since most papers in explainable segmentation do not make this distinction, this might avoid unnecessary confusion when discussing their contents. It should also be noted that interpretability and ease of understanding vary according to the specific audience, whether it be the general public or a more specialized group with specific training, such as radiologists.


\* Corresponding author.

E-mail address: [rokas.gipiskis@mif.vu.lt](mailto:rokas.gipiskis@mif.vu.lt) (R. Gipiškis).

Peer review under responsibility of The Korean Institute of Communications and Information Sciences (KICS).

<https://doi.org/10.1016/j.ict.2024.09.008>

2405-9595/© 2024 The Author(s). Published by Elsevier B.V. on behalf of The Korean Institute of Communications and Information Sciences. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).



**Fig. 1.** Publications with “explainable AI”, “interpretable AI”, and “AI regulation” as keywords. Publication data gathered from [app.dimensions.ai](https://app.dimensions.ai).

XAI is not a new development, particularly in rule-based expert systems [4,5] and machine learning (ML) [6], but it has experienced unprecedented growth ever since the revived interest [7] in neural networks. This growth correlates with the increasing interest in DL and is further driven by (1) the need for trustworthy models due to widely expanding industrial deployments; (2) bureaucratic and top-down political emphasis on AI regulation; and (3) concerns within the ML safety community [8] about the general trajectory of AI development in the short and long runs. AI deployment is increasing across different sectors, and is significant both in terms of its size and impact. According to AI Index Report 2023 [9], the proportion of companies adopting AI has more than doubled from 2017 to 2022. In 2022, the medical and healthcare sectors have attracted the most investment, with a total of 6.1 billion dollars [9]. IBM Global AI Adoption Index 2023 [10], conducted by Morning Consult on behalf of IBM, indicates that about 42% of their surveyed (>1000 employees) enterprise-scale companies reported actively deploying AI, and an additional 40% exploring and experimenting with AI, out of which 59% reported an acceleration in their rollout or investments. Even with rapid deployment, critical high-impact sectors have to move at a slower pace. One could expect even more healthcare-related applications and clinical deployments if AI methods were more interpretable. To a large extent, this applies to other industries as well. According to the same IBM report, most of the surveyed IT professionals (83% among companies already exploring or deploying AI) stated that it is important to their business to explain how their AI reached the decision. Another accelerating trend is that of AI regulation (Fig. 1). The recent survey [11] indicates that 81% of respondents ( $N > 6000$ ) expect some form of external AI regulation, with 57%–66% of respondents reporting that they would be more willing to use AI systems if trustworthiness-assuring mechanisms were in place. AI trustworthiness and transparency are further emphasized in regulatory discussions, ranging from the EU’s AI Act [12] to AI executive order [13] in the United States.

XAI in image segmentation is a relatively new field, with the first articles on the subject appearing in the late 2010s [14–16]. Since then, the topic has gained more attention. Semantic image segmentation is an essential task in computer vision,

with applications ranging from autonomous driving [17] to medical image analysis [18]. Its study is further motivated by the rapidly growing remote sensing and video data. Increasing deployments in medical AI are also contributing to the need for explainable segmentation. Both radiologists and surgeons need to know accurate boundaries for the anatomical structures of interest. Precise and reliable segmentation is required when working with most pathologies in different imaging modalities, ranging from magnetic resonance imaging (MRI) to computed tomography (CT).

Segmentation is commonly viewed as a dense prediction task where classification is performed on a pixel level. However, most XAI literature so far has focused on image classification tasks. Nonetheless, a growing number of works address the issue of interpreting semantic segmentation results by either extending classification-based methods or by proposing their own modifications. Two Ph.D. dissertations [19,20] on XAI in image segmentation have been written in the past year. Image segmentation methods have been reviewed in the medical domain [21], however, the focus has been just on the post-hoc techniques. In this work, we provide the first comprehensive survey in the area of explainable semantic segmentation, encompassing different application domains and all XAI method groups, including architecture-based explanations. The surveyed papers were selected using the Google Scholar search engine, prioritizing such keywords composed of key phrases such as “interpretable”, “explainable”, “interpretability”, “explainability”, “explanation”, “XAI”, “explainable AI”, in combination with “semantic segmentation”, “image segmentation”, and “dense prediction”. Additionally, we employed both backward and forward citation searching, looking for interpretable-segmentation-related papers in both the references of selected papers and in other papers that cite them. Throughout the paper, when referring to the date of the publication, we indicate the date of its first online appearance, including the preprint version where applicable.

This survey offers the following contributions:

- A comprehensive review of up-to-date publications, covering both their theoretical contributions and practical applications.
- A taxonomy to distinguish various interpretability techniques based on five subgroups.
- Analysis of the literature based on evaluation metrics, datasets used, and application domains.
- A detailed discussion of open issues and identification of future research directions.

The paper is structured as follows (Fig. 2). Section 2 begins with the general scope of the problem and then provides the background for the fields of XAI and semantic image segmentation. Section 3 reviews the most important taxonomical dichotomies in classification and introduces a method-centered taxonomy for XAI in image segmentation. Section 4 presents illustrative examples of each method group, includes formalizations for gradient-based and perturbation-based methods, and outlines XAI evaluation metrics. Short summaries of each method with their main contributions, grouped by application

Sections I and II	Sections III and IV	Section V	Sections VI and VII
XAI in Computer Vision	Taxonomy	Medicine	Open Issues
Semantic Segmentation	Methods	Industry	Future Directions
Limitations	Metrics	Self and Weakly Supervised	4. Future Developments
1. Motivation and Background	2. Theory	3. Applications	

Fig. 2. Roadmap of this study.

area, are provided in Section 5. Lastly, Section 6 points out future research directions, while Section 7 draws the main conclusions from this study.

## 2. Background

### 2.1. Development of the field of XAI in computer vision

There is a great variety of XAI methods in classification, with new techniques being proposed weekly. Typically, these methods employ some form of feature attribution, indicating the model’s sensitivity or insensitivity to various features, such as certain pixel configurations in the input space. The most popular explainable classification methods, still influential in today’s DL models, fall into gradient-based or perturbation-based categories. Here, we only highlight key developments, particularly focusing on the methods that have influenced interpretable image segmentation. For an accessible introduction to and treatment of explainable classification, we refer the reader to [22]. For a more detailed survey on these topics, [23] is also recommended.

The first gradient-based explainability techniques for classification in convolutional neural networks (CNNs) are proposed in [24]. The initial method generates artificial images that maximize the score for the selected class of interest. The second method, also referred to as vanilla gradient, produces a saliency map that highlights important regions in the input space. This is based on the gradient for the class of interest with respect to the input. The authors also observe that this method can be used for weakly supervised segmentation. This marks the possibility to use XAI tools instrumentally, not just for the sake of explainability. In [25], the influential Grad-CAM technique is introduced. Its calculation is based on the gradient flow into the last convolutional layer. Since Grad-CAM is calculated for intermediate model activations, the resulting explanation needs to be upsampled. This upsampling process might negatively impact the quality of pixel-level explanations [26]. Similar to [24], the Grad-CAM technique also demonstrates the potential for instrumental use in weakly supervised localization.

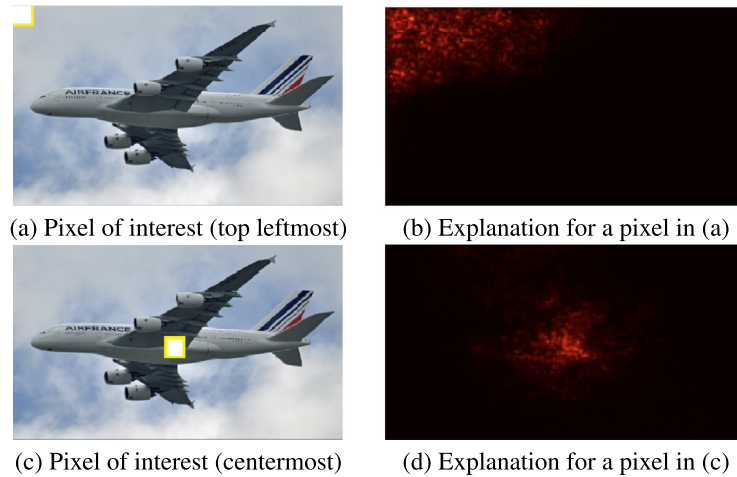
Another area of explainable classification methods encompasses occlusion or perturbation-based techniques, such as occlusion sensitivity measurements [27], LIME [28], SHAP [29], and RISE [30]. In [27], occlusion sensitivity is introduced. It proposes systematically occluding the input image with a smaller gray filter and measuring the effect on the model’s output. The likelihood of the model classifying the image as belonging to the actual class should decrease when the object of that class is occluded in the input space. Other noteworthy methods in explainable classification have focused on optimization. Activation maximization, previously proposed

in [31], initially focused on Restricted Boltzmann Machines, a type of unsupervised models. In [24], it has been specifically implemented in supervised classification models. In [32], this technique was further popularized by demonstrating the results across different network layers. Unlike the previously discussed XAI techniques, this type of explanation method can be described as global because the generated image does not depend on a particular input image but rather on the model’s internal weights.

### 2.2. Specifics of semantic segmentation

Most of the literature on interpretable computer vision focuses on classification. However, DL-based semantic segmentation techniques have achieved significant results. Classical encoder–decoder models such as U-Net [33] or SegNet [34] as well as their modifications, have been deployed in various fields. Vision transformer-based segmentation architectures have also been proposed [35]. There have even been attempts to combine these two approaches [36]. During semantic segmentation, class labels are assigned to each pixel, and the output is typically the same resolution as the input image. Modern segmentation models can be composed of millions of parameters, making their interpretation difficult and often resulting in their description as “black boxes”.

Interpretability in semantic image segmentation is a challenging area of study. On one hand, it can be viewed as an extension of a relatively intuitive interpretable classification. However, it requires combining the relative influence of each classified pixel of interest. On the other hand, interpreting its own explanations is not so straightforward or intuitive. One problem with interpretability methods, not limited to semantic segmentation, is that we generally lack ground truths for the explanations. Furthermore, it is uncertain what the ideal explanation should look like or whether one interpretable instance can be limited to a single explanation. However, in the case of classification, we can at least have some candidates for good explanations. Based on this, qualitative human-based studies [37] can be conducted. Conducting a similar study for semantic segmentation is more complex, as it is less clear what constitutes good explanation candidates: should the interpretability saliency map focus on the entire area of the class of interest or just its boundaries? Can there be instances where the most salient features are outside the class area? What if the segmentation area is correct, but the attributed class is not? Moreover, semantic image segmentation is notorious for inter-observer variability, especially in manual delineations in medical images. One way to demonstrate the usefulness of explainable segmentation is to detect instances where the segmentation of one semantic class appears heavily dependent on the presence of different class pixels, whether nearby or



**Fig. 3.** Explanation for single pixels: The selected pixels are shown on the left, with their corresponding gradient-based explanations on the right. The input images come from the COCO dataset [38].

otherwise. In [16], such a case is demonstrated when the U-Net detects the sky primarily due to the nearby trees, which belong to the ‘Nature’ class. Interpretable semantic segmentation techniques prove most useful when the segmentation is incorrect.

Since we can frame the segmentation task in terms of classification, it is relatively easy to apply explainable classification methods to it, focusing on a single pixel as seen in Fig. 3. For instance, a gradient for the selected output pixel of a chosen class can be calculated with respect to the entire input image. However, an explanation map for the classification of a single pixel is not particularly useful. It is less accessible to the human interpreter, as evaluating thousands of different explanations for just a single class in a single image would be required. Therefore, we need to consider the effects of a larger number of pixels. Most popular explainable segmentation techniques operate under the underlying assumption of pixel importance. This assumption is particularly relevant to perturbation-based methods (Fig. 4), where introducing noise to important pixels would degrade a model’s performance more significantly than adding it to less critical pixels. To explain the whole image (i.e., all pixels) instead of just a single pixel, most explainable segmentation techniques must visualize the relative contributions of all pixels simultaneously. Otherwise, the analysis of separate single-pixel-based explanation maps would be too tedious. The most popular way to do it involves using logit values, unnormalized probabilities before the Softmax layer, typically used in classification. This could be achieved by summing up the logits of the class of interest for the pixels of interest, for instance. This new scalar value can then be used when generating a single explanation for the entire image, just like in the case of a single pixel.

### 2.3. Limitations


Feature attribution and saliency-based XAI methods in particular have faced criticism [39–41]. Although these criticisms have solely focused on explainable classification, they deserve

a thorough examination as they could also extend to segmentation. Some of the XAI methods act as regular edge detectors, independently from the underlying model and training dataset. This independence is troubling because a local post-hoc XAI method should explain a specific model’s prediction for a particular data point. In [42], limitations of feature attribution methods such as SHAP and integrated gradients are emphasized both theoretically and empirically, showing that they cannot reliably infer counterfactual model behavior. The authors observe that the analyzed attribution methods resemble random guessing in tasks like algorithmic recourse and spurious feature identification. Similar experimental results are observed with gradients, SmoothGrad [43], and LIME.

Attribution methods have also been criticized for confirmation bias [44]. An appealing but incorrect explanation might be judged more favorably than a more realistic one. A better understanding of the goals of an idealized attribution method is needed to develop improved quantitative tools for XAI evaluation [44]. In [45], the limitations of post-hoc explanations are investigated. The authors question their effectiveness in detecting unknown (to the user at test time) spurious correlations. These inefficiencies are detected in three types of post-hoc explanations: feature attribution, concept activation, and training point ranking. However, the authors acknowledge that these three classes do not fully cover all post-hoc explanation methods. Other methods have been criticized for their weak or untrustworthy causal relationships. In [46], saliency maps are criticized for their frequent unfalsifiability and high subjectivity. The study also highlights their causal unreliability in reflecting semantic concepts and agent behavior in reinforcement learning environments. In [47], it is argued that feature attribution techniques are not more effective than showing the nearest training-set data point when tested on humans. The limitations of attribution methods in cases of non-visible artifacts [48] have also been investigated.

Despite the critical studies on explainable classification and their potential extensions to segmentation, the widespread





**Fig. 4.** Explanation for multiple pixels: The application of a perturbation-based method to the COCO dataset [38] using summed up logit values.

prevalence of image segmentation requires investigating different explainability tools and their working mechanisms. Although some studies point out the limitations of these techniques, better alternatives have yet to be developed. As observed in [49], the development of interpretability methods is dialectical: a new method is introduced, its failure modes are identified, and as a result, a new method is proposed, with the ongoing aim of making them more reliable. Current methods have much room for improvement, especially considering that the entire field is in the early stages of development. The above criticisms can serve as sanity checks for XAI methods. Despite the limitations, some techniques, such as gradients and Grad-CAM in the case of [39], do pass certain sanity checks. Even some critical literature [49] agrees that certain explainability techniques can be useful for exploratory use cases. To our knowledge, the specifics of XAI limitations in image segmentation have not yet been explored.

### 3. Taxonomy

Different XAI taxonomies have been introduced in classification, both with respect to specific subgroups of interpretability methods [50,51], and with respect to more abstract conceptual terms [52]. Even meta-reviews of various existing taxonomies have been proposed [53,54]. Since image segmentation can be seen as an extension of classification, many taxonomy-related aspects can be validly transferred from research in explainable classification. In most taxonomies, a particularly important role is played by three dichotomies: post-hoc vs ad-hoc (sometimes also referred to as inherent interpretability), model-specific vs model-agnostic, and local vs global explanations. In this section, we briefly cover these high-level dichotomies and then present a low-level taxonomy based on the surveyed papers, categorizing various interpretability techniques into five subgroups based on how the explanation is generated.

#### 3.1. Scope: Local vs. Global

The first prevalent dichotomy distinguishes between local and global explanations. Here, locality refers to the use of a single input image with respect to which the explanation is given. A global explanation, on the other hand, would aim to explain the model's behavior across a range of different images, not limiting itself to just one. According to meta-surveys [53,54], the local–global dichotomy is prevalent in numerous XAI taxonomies. This distinction is essential in explainable segmentation as well, with most methods falling under local explanations.

#### 3.2. Method and its timing: Post-hoc vs. Ad-hoc

The distinction between post-hoc and ad-hoc explanations highlights that one can either apply XAI techniques to an already-trained model without any interference or apply them during and as part of the training process. Sometimes, these explanations are also described as passive and active approaches [23]. Under this definition, active approaches require modifications to the network or the training process. Such changes influence both the model's performance in terms of evaluative metrics and its interpretability. Therefore, an accuracy-interpretability trade-off cannot be avoided in ad-hoc XAI methods, but it is avoided in the case of post-hoc applications.

This widely accepted dichotomy can nonetheless be slightly misleading, as both terms can be meant to emphasize different distinctive criteria. Post-hoc can be understood as referring to the fact that XAI techniques are applied after the training, hence “post”. Naturally, it would seem that ad-hoc should be understood as referring to XAI techniques that are applied during training. However, sometimes, as a direct opposition to “post-hoc”, terms like ‘inherent interpretability’ [2] or “self-explainability” [55] are used, pointing to an entirely different aspect: the architecture or the type of XAI method. In some cases, such interpretation could allow for XAI methods that are both inherently interpretable and post-hoc [22], which might cause confusion.

#### 3.3. Range: Model-specific vs. Model-agnostic

The third distinction evaluates the flexibility of a given XAI technique in its application to different model architectures. Model-specific XAI methods heavily depend on the underlying model architecture, whereas model-agnostic methods are more universal in their compatibility with various models, and can be applied to different architectures without further modifications. The interpretation of inherently interpretable models is always model-specific [22].

#### 3.4. XAI taxonomy for image segmentation

Multiple compatible taxonomies are possible depending on the level of abstraction in which we are interested. In [50], XAI methods in ML are divided into transparent models and post-hoc explainability, which is then further divided into model-specific and model-agnostic categories. In [51], interpretation methods are divided into post-hoc interpretability analysis and ad-hoc interpretable modeling. In [56], a

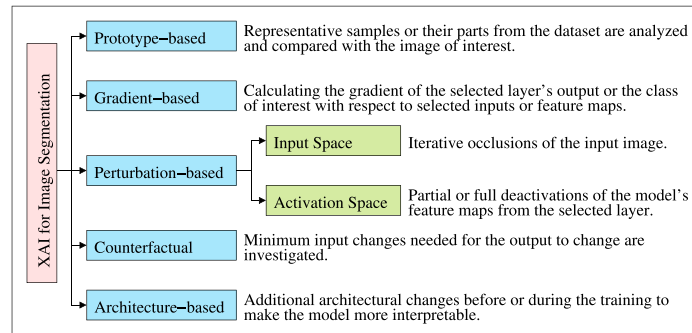


Fig. 5. Method-centered taxonomy for explainable image segmentation.

higher-level taxonomy distinguishes between structural analysis, behavioral analysis, and explainability by design. In [57], a preliminary taxonomy of human subject evaluation in XAI is introduced, which might be particularly useful when using qualitative evaluations of XAI. Based on the analysis of XAI taxonomies in classification, we observe that they could also be applied to image segmentation. However, no specific framework has been introduced to address the ever-growing field of interpretable segmentation. We hope that a more detailed demarcation will be useful in navigating across different types of techniques. In our survey, we propose a taxonomy (Fig. 5) that is based on the reviewed literature in explainable image segmentation.

The proposed method-centered taxonomy includes five method families: prototype-based, gradient-based, perturbation-based, counterfactual methods, and architecture-based techniques. Based on the previously discussed high-level taxonomies commonly presented in other surveys, most of the techniques fall under local and post-hoc explainability categories. Exceptions are primarily found in the architecture-based category, which includes examples of ad-hoc, global, and model-agnostic explanations. Prototype-based methods employ representative samples or their parts from the dataset to analyze and compare with the input image. Gradient-based methods involve calculating the gradient of the output of a selected layer or the class of interest with respect to selected inputs or feature maps. Perturbation-based methods can be divided into two groups based on the perturbed space. Input space perturbations are iterative occlusions of the input image. Typically, they are based on a sliding filter, but different types of noise can also be introduced. Explanations are based on their effect on the model's outcome. Activation space perturbations involve partial or full deactivations of the model's feature maps from the selected layer. Once again, explanations are based on their effect on the model's outcome. Counterfactual methods employ the minimum input changes needed for the output to change. Finally, architecture-based techniques involve making additional architectural changes either before or during training to enhance interpretability. Section 4 presents a more detailed analysis of each method group.

#### 4. XAI for image segmentation

In this section, we review the main methods representative of each subgroup in the taxonomy, as well as the metrics for explainable image segmentation.

#### 4.1. Methods

##### 4.1.1. Prototype-based methods

Prototype-based models [58] utilize typical representatives from the dataset, usually selected from the training set. These methods emphasize the intuitiveness of the provided explanations, presenting them in an easily understandable form of naturally occurring objects. Such features can be easily distinguished and discriminated by end users. Meanwhile, prototypical parts refer to specific regions within representative prototypes, also known as exemplars. In contrast to a prototype, a criticism is a data instance that is not well represented by the prototypes [22]. In terms of architecture, typical prototype-based methods require the insertion of a prototype layer into the segmentation model. Therefore, depending on the taxonomy, prototype-based methods could also be viewed as self-explainable and part of the architecture-based methods. However, due to their frequent mentions in the related classification literature under the same subgroup label, we opt to treat them as a separate group.

Although prototypical methods are prevalent in classification [59–61], their extensions for segmentation are few. Typically, prototype layer (Fig. 6) is a key component in prototype-based methods for both classification [59–61] and segmentation [55,62]. Within a prototypical layer, different classes are represented by predefined or learned prototypes. In [55], a ProtoSeg model is proposed. The authors introduce a diversity loss function based on Jeffrey's divergence [63] to increase the prototype variability for each class. Better results are observed when the diversity loss component is introduced. The authors attribute this to the higher informativeness of a more diverse set of prototypes that leads to a better generalization. We think that this could be related to the diversity hypothesis [64], first introduced in the context of reinforcement learning, and could be explored further. The experiments are performed using Pascal VOC 2012 [65], Cityscapes [66], and EM Segmentation Challenge [67] datasets. DeepLab [68] model is used as the backbone. In [62], a prototype-based method is used in combination with if-then rules for the interpretable segmentation of Earth observation data. The proposed approach is the extension of xDNN [69], and uses mini-batch K-mean [70] clustering. For the feature extraction part, the U-Net architecture is used. The experiments are performed using the Worldfloods [71] dataset.

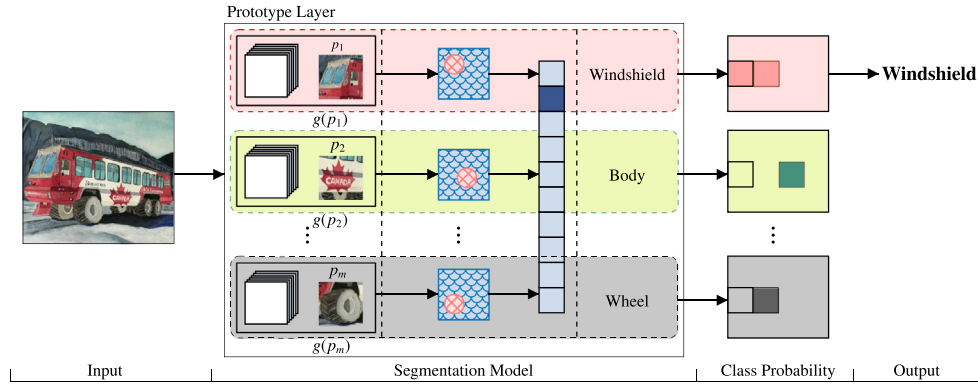


Fig. 6. A framework for prototype-based methods.




Fig. 7. A framework for counterfactual methods.

#### 4.1.2. Counterfactual explanations

Counterfactual or contrastive explanations investigate the minimum input changes needed for the output to change. Unconditional counterfactual explanations were first introduced in [72]. This explainability subfield is related to adversarial attacks. Counterfactual images are similar to the original, yet are able to change the model's output. Counterfactual explanations can also be viewed as closely linked to perturbation-based explanations, which will be discussed in the next subsection. Counterfactual XAI techniques frequently fall into local post-hoc category [73]. After the initial segmentation model, counterfactual-based interpretability method typically employs additional networks for counterfactual generation. In our pipeline (Fig. 7), this is depicted by additional encoder and decoder networks.

Generator-based counterfactual explanations are investigated in [74]. OCTET, a generative approach, produces object-aware counterfactual explanations for complex scenes. Counterfactual changes to the image focus on road markings, such as changing the solid line into a dashed one, or the positions of cars by cropping and extending the relevant regions of the input image. The models are trained on the BDD100k [75] and BDD-OIA [76] datasets. Additional information can be found in the supplementary material [77]. In [78], segmentation results are qualitatively compared using counterfactual images. The experiments are performed on Kvasir-seg [79] and Kvasir-instrument [80] datasets. Counterfactual explanations are generated using the segmented area of interest, which is then replaced with the average pixel value of the rest of the image. In [81], counterfactual explanations are generated for complex scenes while preserving the semantic structure. The proposed method uses semantic-to-real image synthesis. Here, a noticeable contrast can be drawn between this approach and perturbation-based methods. In the latter, perturbations applied

to the input space fail to produce semantically meaningful image regions. While perturbation-based techniques are related to counterfactual explanations, they form a distinct class, even though counterfactual explanations may also include perturbations.

#### 4.1.3. Perturbation-based methods

Perturbation-based methods typically employ occlusions in the input (Fig. 8) or activation (Fig. 9) space, and then measure their influence on the model's output for a selected class. Here occlusions or perturbations could be understood as uninformative regions, transforming the input or its internal feature maps. Pixels in the occlusion filter can be set to 0, as seen in Fig. 8, or any other arbitrary value. Such sliding filter would occlude different regions of the input space. Its size and stride parameters are specified beforehand. Gaussian or any other random noise can also be used for these purposes. Multiple perturbative iterations are required for the generation of an explanation map. During each inference, the score is calculated, measuring the perturbation's effect on model's performance. This can be done by taking the difference between the score for the original image and that of its perturbed version. Such a score can be based on an evaluative metric or pre-Softmax prediction values. Since the same input image has to undergo multiple transformations, each requiring a separate forward pass through the model, perturbation-based XAI methods are considered computationally expensive. Another limitation is that perturbation-based input modifications can sometimes produce images that lie outside the original training distribution. These modified images may not resemble any data the model has previously encountered, which can lead to explanations that are not representative of the model's typical decision-making process. This, in turn, may generate




Fig. 8. A framework for perturbation-based methods for the input space.




Fig. 9. A framework for perturbation-based methods for the activation space.

misleading or irrelevant insights into how the model operates on actual, in-distribution data.

Following [82], given an RGB image  $x$  of dimensions  $N \times M \times 3$ , the pre-Softmax prediction score for pixel  $x_{ij}$  for class  $c$  is  $l^c(x_{ij})$ . When  $x_{ij}$  is classified as  $c$ , the sum of these scores is:

$$L^c(x) = \sum_{i,j} [\hat{c}_{ij} = c] l^c(x_{ij}), \quad (1)$$

where  $\hat{c}_{ij}$  is the predicted class for  $x_{ij}$ .

Following [83],  $L^c(x)$  is used to compute the importance weight  $w_k^c$  of each activation map  $k$ :

$$w_k^c = \frac{L^c(x) - L_k^c(x)}{L^c(x)}, \quad (2)$$

where  $L_k^c(x)$  is the sum of pre-Softmax scores for  $c$  after deactivating  $k$ . These calculated importance scores are then used in a linear combination of feature maps to generate a perturbation-based explanation.

The authors of [15] propose the first XAI solution for extending saliency techniques beyond classification. Their perturbation-based method is introduced for the detection of contextual biases. The experiments are performed using a synthetic toy dataset based on MNIST [84] as well as the Cityscapes [66] dataset. In [85], a hybrid SegNBDT approach is introduced, combining both decision trees and neural networks. This method falls under both the perturbation-based and self-explainable model categories. For the experimental part, Pascal Context [86], Cityscapes [66], and Look Into Person [87] datasets are used. In [88], SHAP and RISE techniques are applied to image segmentation. SHAP is a popular post-hoc interpretability method, and the proposed approach is based on Kernel SHAP [29]. The experiments are performed on Synthetic Aperture Radar images from the unspecified dataset for oil slick detection at the sea surface and the Cityscapes [66] dataset. In [26], a perturbation-based occlusion sensitivity approach is used to measure the performance of the proposed interpretable semantic segmentation approach. Compared to occlusion sensitivity and Grad-CAM, their method

achieves orders of magnitude lower inference time. However, it requires training an additional interpretability model. In [89], following [27], different types of input occlusions are investigated for applications in semantic segmentation. The paper discusses how occlusion filter sizes and colors can affect the generated explanations. It is observed that, compared to image classification, input occlusions in segmentation models do not generate as much variance in the evaluation metric scores. For the experimental investigation, COCO [38] dataset is used. The proposed method is evaluated qualitatively, with select images also compared quantitatively using deletion curves.

Perturbations are not limited to the input space. For instance, Ablation-CAM [83] is a gradient-free method that systematically deactivates feature maps in a selected layer. In [82], Ablation-CAM is extended to semantic segmentation. It is a gradient-free interpretability technique based on ablating or perturbing activation maps. The experiments are performed on a private industrial dataset for fruit-cutting machines as well as on COCO [38] dataset.

#### 4.1.4. Gradient-based methods

Gradient-based methods (Fig. 10) typically use gradients of the outputs from later layers with respect to the input features. These techniques are less computationally expensive compared to perturbation-based techniques because only a single backward pass is required. Perturbation techniques, on the other hand, require a separate forward pass for each perturbed image, increasing computational costs with each inference. However, gradient-based saliency maps can generate more noise.

Following [90], given an RGB image  $x$  of dimensions  $N \times M \times 3$ , and a set of class labels  $\{1, 2, \dots, C\}$ , where  $C$  denotes the total number of classes, we define:

$$g(x) = (g_1(x_{ij}), \dots, g_C(x_{ij})) \in \mathbb{R}^{N \times M \times C}, \quad (3)$$

where  $g_c(x_{ij})$  denotes the pre-Softmax prediction score for class  $c$  at pixel  $x_{ij}$ .

The sum of these pre-Softmax scores for class  $c$  is then defined as:

$$g_{c,A}(x) = \sum_{i,j \in A} g_c(x_{ij}), \quad (4)$$

where  $A$  is a set of pixel indices of interest.

From this, the gradient-based explanation with respect to  $c$  is derived as:

$$G_A(x, c) = \frac{\partial g_{c,A}(x)}{\partial x}. \quad (5)$$

In [16], Seg-Grad-CAM is proposed as the extension of Grad-CAM [25]. It is one of the best known explainability techniques in image segmentation. Just like in the case of regular Grad-CAM, the generated saliency is based on the weighted sum of the selected feature maps. Its application is demonstrated on a U-Net model, trained on the Cityscapes [66] dataset. In [91], the same method is applied for automatic rock joint trace mapping. The original Grad-CAM technique for classification, together with simple gradients, passes the previously discussed sanity checks, evaluating the reliability



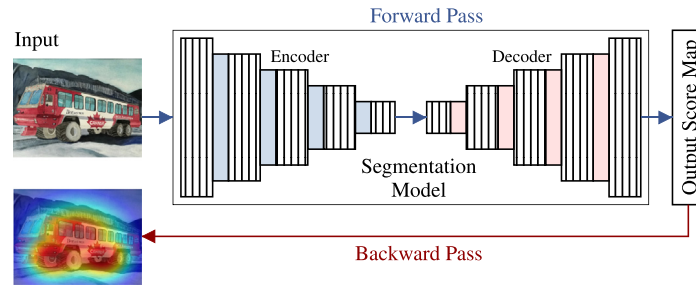


Fig. 10. A framework for gradient-based methods.




Fig. 11. A framework for architecture-based methods.

of XAI technique. In [92], Seg-XRes-CAM is introduced. The authors criticize Seg-Grad-CAM [16] for not utilizing spatial information when generating saliency maps for a region of the segmentation map. The proposed approach draws inspiration from HiResCAM [93], a modification of the original Grad-CAM [25]. Subsequently, [94] adapts five CAM-based XAI methods from classification to the segmentation of high-resolution satellite images. Among the proposed extensions are Seg-Grad-CAM++, Seg-XGrad-CAM, Seg-Score-CAM, and Seg-Eigen-CAM. Just like in [82], Ablation-CAM, a gradient-free method, is also extended for segmentation. Besides using the drop in the segmentation score to measure their methods' performance, the authors also propose entropy-based XAI evaluation metric. The implemented methods are tested on a WHU [95] building dataset. In [96], an interpretability and visualization toolbox is proposed for classification and segmentation networks. It includes several XAI extensions specifically for image segmentation. Among them are Guided Grad-CAM and segmented score mapping, extended from [97].

#### 4.1.5. Architecture-based methods

This subgroup of methods introduces additional architectural changes (Fig. 11) that aim to make the models more interpretable. Instead of relying on post-hoc techniques that are added on top of the already trained models, these methods are typically employed as part of the training process. This class of XAI methods is sometimes described as interpretable by design, inherently interpretable, or interpretability as part of the architecture. In this case, interpretability is inherently linked to the specific model architecture or design, and is not easily transferable to other architectures.

One such example is chimeric U-Net with an invertible decoder [117]. This approach introduces architectural constraints for the sake of explainability. The authors claim that it can

achieve both local and global explainability. In [161], both supervised and unsupervised techniques of Semantic Bottlenecks (SB) are introduced for better inspectability of intermediate layers. This approach is proposed as an addition to the pre-trained networks. Unsupervised SBs are identified as offering greater inspectability compared to their supervised counterparts. The experiments are primarily performed on street scene segmentation images from the Cityscapes dataset. The results are also compared using two other datasets: Broden [162] and Cityscapes-Parts, a derivative of Cityscapes. In [104], a framework for symbolic semantic segmentation is proposed. This work is at the intersection of image segmentation and emergent language models. The authors apply their research to medical images, specifically brain tumor scans. Emergent Language model with a Sender and a Receiver is utilized for interpretable segmentation. The Sender is an agent responsible for generating a symbolic sentence based on the information from the high model layer, while the Receiver cogenerates the segmentation mask after receiving symbolic sentences. Symbolic U-Net is trained on the Cancer Imaging Archive (TCGA) dataset<sup>1</sup> and used for providing inputs to the Sender network.

#### 4.2. Metrics

XAI techniques are used in addition to standard evaluation metrics due to their limitations. However, to evaluate the performance of these techniques, they also need to be measured. We can distinguish between qualitative and quantitative assessment methods. Qualitative evaluation commonly refers to user-based evaluation and, based on the surveyed papers (Tables 1 and 2), is the more prevalent of the two. To quantify subjective user results, various questionnaires have been proposed [179], such as the explanation goodness checklist, explanation satisfaction scale, trust scales, and the ease of understanding when comparing different explainability techniques [180]. These methods still require polling multiple subjects, although, when surveying experts, in practice their number is limited to 2–5 [180]. This way, quantification still takes place, but it is based on subject-dependent evaluation. Since questionnaire studies require additional resources, most of the papers using qualitative evaluation only provide visual comparisons between different XAI techniques, leaving qualitative evaluation to the reader's eye.

<sup>1</sup> <https://wiki.cancerimagingarchive.net/pages/viewpage.action?pageId=5309188>.

**Table 1**

Explainable image segmentation in medicine.

Field	Imaging modality	Objects of interest	Datasets	Metric	Year	Ref.
G	IMG <sup>a</sup>	Colorectal polyps	EndoScene [98]	▷	2018	[14]
O	CT	Liver tumors	LiTS [99]	▶	2019	[100]
C	CMRI	Ventricular volumes	SUN09 [101], AC17 [102]	▷	2020	[103]
O	MRI	Brain tumors	TCGA	▷	2020	[104]
D	MRI/IMG	Skin lesions, multi-organ (incl. the fetal brain and the placenta)	ISIC2018 [105] and a private fetal MRI dataset	▷	2020	[106]
P	CT	Pancreatic region	Medical segmentation decathlon	▶	2021	[26]
C	CMRI	Ventricles, myocardium	Cardiac MRI dataset [102]	▷	2021	[107]
G	IMG	Polyps, med. instruments	Kvasir-SEG [79], Kvasir-Instrument [80]	▷	2021	[108]
O	MRI	Brain tumors	BraTS2018 [109]	▶	2021	[110]
FM	NIR	Iris	Private test dataset, post-mortem iris datasets, collected by [111]	▷	2022	[112]
V	CT/MRI/IMG	Skin lesions, abdomen multi-organ, brain tumors	HAM10000 [113], CHAOS 2019 [114], BraTS 2020 [109]	▷	2022	[115]
V	MRI	Brain tumors, human knees	BraTS 2017 [109], OAI ZIB [116]	▷	2022	[117]
O	MRI	Brain tumors	BraTS 2019 [109], BraTS 2021 [109]	▷	2022	[118]
N	MRA	Brain vessels	Private	▶	2022	[119]
H	IMG	Liver	Simulated dataset (Test Set 4) [120]	▷	2022	[121]
O	US/MG	Breast tumors	Private LE/DES datasets, and BUSI [122]	▷	2023	[123]
G	CT/IMG	Colorectal polyps, lung cancer	EndoScene [98], LIDC-IDRI [124]	▶	2023	[125]
O	CT/MRI	Prostate cancer	3D pelvis dataset [126]	▷	2023	[127]
G	CT	Abdominal organs	Synapse multi-organ CT dataset [128]	▷	2023	[92]
O	BUS	Breast tumors	BUSI [122], BUSIS [129], HMSS [130]	▷	2023	[131]
O	CT/PET	Non-small cell lung cancer, whole-body	NSCLC, AutoPET [132]	▶	2023	[133]
D	IMG	Skin lesions	ISIC2018 [105]	▷	2023	[134]
D	IMG	Melanoma	ISIC 2018 [105], ISIC 2019 [113]	▷	2023	[135]
O	MRI/IMG	Prostate tumors, optic disc and cup	Prostate <sup>b</sup> and fundus <sup>c</sup> datasets	▷	2023	[136]
O	X-ray	Breast tumors	INbreast [137]	▶	2023	[138]
O	WSI	Head and neck tumors	Private	▷	2023	[139]
A	IR	Feet	ThermalFeet	▷	2023	[140]
V	CT/MRI/IMG	Brain tumors	BraTS 2018 [109], BraTS 2019 [109], BraTS 2020 [109], ISIC 2017	▷	2023	[141]
P	CT	Pancreas	Pancreas segmentation dataset [142]	▷	2023	[143]
Op	OCT	Retinal layers, glaucoma, diabetic macular edema	NR206, glaucoma dataset [144], DME dataset [145]	▷	2023	[146]
V	CT/MRI	Prostate, left ventricle, right ventricle, myocardium	NCI-ISBI 2013 [147], I2CVB [148], PROMISE12 [149]; MSCMR [150], EMIDEC [151], ACDC [102], MMWHS [152], CASDC 2013 [153]	▷	2023	[154]
Op	OCT	Retinal layers, glaucoma, diabetic macular edema	Vis-105H, glaucoma dataset [144], DME dataset [145]	▷	2024	[155]
V	CMRI/CT	Left atrium, thoracic organs	Atrium dataset [142], SegTHOR [156]	▷	2024	[157]

A: Anesthesiology, C: Cardiology, D: Dermatology, FM: Forensic Medicine, G: Gastroenterology, N: Neurology,

O: Oncology, Op: Ophthalmology, P: Pancreatology, and V: Various.

▷: Qualitative XAI evaluation.

▶: Quantitative XAI evaluation.

<sup>a</sup> IMG: general-purpose digital image formats, such as JPEG.<sup>b</sup> Prostate datasets: RUNMC [147], BMC [147], HCRUDB [148], UCL [149], BIDMC [149], and HK [149].<sup>c</sup> Fundus datasets: DRISHTI-GS [158], RIM-ONE-r3 [159], and REFUGE [160].

Quantitative evaluation does not involve human subjects and can be more easily applied when comparing different interpretability methods. Infidelity and sensitivity [181] are the only two metrics that, as of 2024, are implemented in the Captum [182] interpretability library for PyTorch. Deletion and insertion metrics [30] are another type of quantitative evaluation, based on measuring the Area under the Curve (AUC), generated after gradually deleting or inserting the most important pixels in the input space. A wider use of quantitative metrics could allow for a more standardized comparison across different models and explainability techniques, eliminating the element of human subjectivity. They could

also be valuable when human studies are impractical due to resource constraints. This scalability is particularly relevant in industrial and research settings where numerous models need to be evaluated quickly and consistently. However, for some XAI methods, such as counterfactual explanations, it might be difficult to evaluate the usefulness of the explanation quantitatively. In the case of counterfactual explanations, we can measure whether the generated images are realistic and how closely they resemble the query images, but for a more thorough evaluation of the explanation itself, user studies [74] might be required.

**Table 2**  
Explainable image segmentation in industry.

Category	Domain	Datasets	Metric	Year	Ref.
Remote sensing	Building detection	IAIL [163]	▷	2019	[164]
Scene understanding	Autonomous driving	SYNTHIA [165], A2D2 [166]	▷	2021	[167]
Scene understanding	Pedestrian environments	PASCAL VOC 2012 [65], ADE20K [168], Cityscapes [66]	NA <sup>a</sup>	2021	[169]
Scene understanding	Autonomous driving	KITTI [170]	▷	2022	[171]
Environmental monitoring	Flood detection	Worldfloods [71]	▷	2022	[62]
Scene understanding/Biometrics	Driving scenes/Face recognition	BDD100k [75], CelebAMask-HQ [172], CelebA [173]	▷	2022	[81]
Monitoring/Scene understanding	Drones/Food processing	ICG drone dataset, private dataset	▷	2023	[90]
Monitoring/General applications	Food processing	COCO [38], private dataset	▷	2023	[82]
Biometrics	Facial emotions	Face recognition dataset [174]	▷	2023	[175]
Monitoring	Cracks in infrastructure	CrackInfra [176]	▷	2023	[176]
General applications	Common objects	COCO [38]	▶	2023	[89]
Scene understanding/General applications	Street scenes/Common objects	Pascal VOC 2012 [65], Cityscapes [66]	▷	2023	[55]
Scene understanding	Driving scenes	BDD100k [75], BDD-OIA [76]	▷	2023	[74]
Scene understanding/General applications	Street scenes/Common objects	Cityscapes [66], Pascal VOC [65], COCO [38]	▶	2023	[177]
General applications	Common objects	COCO [38]	▷	2023	[92]
General applications	Common objects	Pascal VOC [65]	▶	2023	[125]
Scene understanding/Remote sensing	Street scenes/Building detection	Cityscapes [66], WHU [95]	▶	2023	[178]

▷: Qualitative XAI evaluation.

▶: Quantitative XAI evaluation.

<sup>a</sup> The application focuses on introducing explainability to segmentation evaluation, rather than evaluating explainability techniques.

In [183], a psychophysics study ( $N = 1150$ ) is conducted to evaluate the performance of six explainable attribution methods on different neural network architectures. Shortcomings in the methods are detected when using them to explain failure cases. Comparative quantitative rankings of different saliency techniques can also be inaccurate. In [184], inspired by [39], sanity checks for saliency maps are investigated. The authors perform checks for inter-rater reliability, inter-method reliability, and internal consistency, and determine that the current saliency metrics are unreliable. It is observed that these metrics exhibit high variance and are sensitive to implementation details.

## 5. Applications

This section presents concrete XAI applications in medical and industrial domains. We also discuss other use cases, with a primary focus on industry-related monitoring domains, such as remote sensing, environmental observation, and biometrics. Additionally, the potential uses of XAI for self-supervised image segmentation are reviewed.

### 5.1. Medical applications

Most applications in explainable image segmentation have been investigated in the medical domain, using datasets from various medical fields (Table 1), ranging from cardiology to oncology. Proposed XAI solutions and applications are employed for diagnosing, monitoring, and other clinical tasks. In some cases, there might be unavoidable overlaps between medical fields. For instance, overlaps occur at the intersection of oncology and histopathology when discussing microscopic tumor images, or between oncology and dermatology when

considering melanoma [135]. Such overlaps can also arise from using multiple datasets, each associated with a different medical field. In these instances, we specify the relevant details in the method description.

#### 5.1.1. Dermatology

Dermatology-centered XAI applications [134,135] focus on skin lesions. Specifically, [135] discusses applications for interpreting melanoma diagnosis results. The proposed pipeline utilizes both classification and segmentation networks. Grad-CAM is employed to generate explainable heatmaps for the classifier, which are then used as inputs in the U-Net network. These heatmaps assist in generating indicator biomarker localization maps. The proposed approach can be used in self-supervised learning. Experiments are performed on the ISIC 2018 [105] and ISIC 2019 [113,185,186] datasets. In [134], a CAM-based explainability metric is proposed and incorporated into the loss function. This metric quantifies the difference between the CAM output and the segmentation ground truth for the targeted class. Both segmentation and explanation losses are considered during the model's training phase. The use of CAM with learnable weights enables a balance between segmentation performance and explainability. The proposed method belongs to the self-explainable XAI category. Similar to [135], the U-Net network is used. The experiments are conducted on the ISIC2018 [105] dataset. In [106], a comprehensive attention-based convolutional neural network is proposed for better interpretability in dermoscopic and fetal MRI images. This approach uses multiple attentions, combining the information about spatial regions, feature channels, and scales. The experiments are performed on ISIC 2018 [105] and a private fetal MRI dataset.

### 5.1.2. Forensic medicine

The applications of explainable segmentation in forensic medicine are limited to iris segmentation. This can be more narrowly referred to as forensic ophthalmology. In [112], the investigation focuses on forensic postmortem iris segmentation. The authors apply a classical technique of Class Activation Mapping (CAM) [187]. The experiments are performed on a private test dataset, and publicly available post-mortem iris datasets, collected by [111].

### 5.1.3. Gastroenterology

XAI applications for endoscopic image segmentation primarily focus on polyps. In [14], the Guided Backpropagation [188] technique is extended to the semantic segmentation of colorectal polyps. Uncertainty in input feature importance is estimated, with higher uncertainty observed in inaccurate predictions. Uncertainty maps are generated using the Monte Carlo Dropout method. The proposed solution is evaluated on the EndoScene [98] dataset. In [108], Layer-wise Relevance Propagation (LRP), a propagation-based explainability method, is applied to the endoscopic image segmentation of gastrointestinal polyps and medical instruments. LRP is specifically applied to the generator component within a Generative Adversarial Network. The generated relevance maps are then qualitatively evaluated. The segmentation models are trained on the Kvasir-SEG [79] and Kvasir-Instrument [80] datasets.

### 5.1.4. Hepatology

In [121], two gradient-based post-hoc explanations, Grad-CAM and Grad-CAM++, are investigated for cross explanation of two DL models, U-Net and the Siamese/Stereo matching network, based on [120]. The experiments are performed on laparoscopic simulated stereo images [120], with a focus on liver segmentation.

### 5.1.5. Oncology

Most of the explainable medical AI applications in image segmentation are in oncology.

**Liver:** A DeepDream-inspired method is proposed in [100] for the segmentation of liver tumors in CT scans, specifically focusing on binary segmentation. The study seeks to understand how human-understandable features influence the segmentation output and defines the network's sensitivity and robustness to these high-level features. High sensitivity indicates the importance of such features, while high network robustness shows its indifference to them. Radiomic features are also analyzed. The experiments are performed on the LiTS<sup>2</sup> [99] challenge dataset. Semantic segmentation in liver CT images is further investigated in [189], where the segmentation output is corrected based on XAI. This approach is categorized as a global surrogate and is model-agnostic. However, its primary purpose is not interpretability but rather the improvement in the initial segmentation by using additional boundary validation and patch segmentation models. The authors of [190] investigate the segmentation of malignant melanoma lesions in 18-fluorodeoxyglucose (<sup>18</sup>F-FDG)

PET/CT modalities, focusing on metastasized tumors. The claim to interpretability is based on the visualization of the model's intermediate outcomes. The overall pipeline involves both segmentation and detection. Volumes of interest (VOI) are visualized for the liver as well as PET-positive regions classified as physiological uptake. This additional information is provided together with the final segmentation masks.

**Brain:** An interpretable SUNet [104] architecture is proposed for the segmentation of brain tumors using The Cancer Imaging Archive (TCGA) dataset. Experimental results and statistical analysis indicate that symbolic sentences can be associated with clinically relevant information, including tissue type, object localization, morphology, tumor histology and genomics data. In [110], 3D visual explanations are investigated for brain tumor segmentation models, using the quantitative deletion curve metric to compare the results with Grad-CAM and Guided Backpropagation [188] techniques. In [115], a region-guided attention mechanism is used for the explainability of dermoscopic, multi-organ abdomen CT, and brain tumor MRI images. The experiments are performed on HAM10000 [113], CHAOS 2019 [114], and BraTS 2020 [109] datasets. Another architecture-based solution is proposed in [117], where the U-Net architecture is modified and applied to two MRI datasets: BraTS 2017 [109] and OAI ZIB [116], respectively focusing on brain tumors and human knees. In [191], Grad-CAM results are compared to brain tumor segmentation results. The overall pipeline includes both classification and segmentation networks, where DenseNet is used for classification and Grad-CAM-based heatmaps are generated for different layers. However, Grad-CAM is not specifically tailored for segmentation but rather used as an explainable classification tool to evaluate segmentation results. In [118], a NeuroXAI framework is introduced, combining seven backpropagation-based explainability techniques, each suitable for both explainable classification and segmentation. Gliomas and their subregions are investigated using 2D and 3D explainable sensitivity maps. A ProtoSeg method is proposed in [141] for interpreting the features of U-Net, presenting a segmentation ability score based on the Dice coefficient between the feature segmentation map and the ground truth. Experiments are performed on five medical datasets, including BraTS for brain tumors, each focusing on different medical fields or affected organs.

**Pelvis:** In [127], a Generative Adversarial Segmentation Evolution (GASE) model is proposed for a multiclass 3D pelvis dataset [126]. The approach is based on adversarial training. Style-CAM is used to learn an explorable manifold. The interpretability part allows visualizing the manifold of learned features, which could be used to explain the training process (i.e. what features are seen by the discriminator during training).

**Breast cancer:** Oncological XAI applications for the segmentation of breast tumors are investigated in [123,131], and [138]. In [123], a multitask network is proposed for both breast cancer classification and segmentation. Its interpretations are based on contribution score maps, which are generated by the information bottleneck. Three datasets are used,

<sup>2</sup> <https://competitions.codalab.org/competitions/17094>.



each focusing on a different imaging modality. In [131], SHAP explainability method is applied to the task of breast cancer detection and segmentation. The experiments are performed on BUSI [122], BUSIS [129], and HMSS [130] datasets. In [138], explainability for mammogram tumor segmentation is investigated with the application of Grad-CAM and occlusion sensitivity, in both cases using Matlab implementations, and activation visualization. Their quantitative evaluation is based on image entropy, which gives additional information about the XAI method's complexity. Pixel-flipping techniques, which are directly related to deletion curves, are also employed. The experiments are performed on INbreast [137] dataset of X-ray images.

**Other:** In [136], Importance Activation Mapping (IAM) is employed as an explainable visualization technique in continual learning. The generated heatmap shows which regions in the input space are activated by model parameters with high-importance weights, associated with the model's memory. This approach is evaluated for the segmentation of prostate cancer. It also has applications in ophthalmology, specifically for segmenting the optic cup and disc. In [139], two CAM-based XAI techniques, Seg-Grad-CAM and High-Resolution CAM (HR-CAM), are applied to histopathological images of head and neck cancer. The explanations generated by both techniques appear to rely on the same features identified by professional pathologists. In [192], a solution based on Cartesian Genetic Programming is used to generate transparent and interpretable image processing pipelines. This method is applied to biomedical image processing, ranging from tissue histopathology to high-resolution microscopy images, and can be characterized as a few-shot learning approach. In [193], a classification-based version of Grad-CAM is used to enhance a U-Net-based segmentation network. The experiments are performed on the 3D-IRCADb-01 [194] dataset, comprised of 3D CT scans of venous phase CT patients. An Xception network generates 2D saliency maps for classification, which are then passed to the U-Net network together with the corresponding input images. This prior information enables more accurate segmentation. In [195], a framework for explainable classification and segmentation is presented. For segmentation, it relies on a feature hierarchy. The experiments are performed on the skin cancer dataset. The Factorizer architecture, introduced in [196], is based on nonnegative matrix factorization (NMF) components, which are argued to be more semantically meaningful compared to CNNs and Transformers. The proposed approach is categorized under architecture-based interpretability methods. The models are implemented for brain tumor and ischemic stroke lesion segmentation datasets. In [119], a framework for explainable semantic segmentation is presented, extending several classification techniques to segmentation. These methods are also applied to 3D models. Infidelity and sensitivity metrics are used, and the experiments are performed on vessel segmentation in human brain images using Time-of-Flight Magnetic Resonance Angiogram. The experimental data [197] is not publicly available. In [133], a new interpretation method is proposed for multi-modal segmentation of tumors in PET and CT scans. It introduces a novel loss function to facilitate

the feature fusion process. The experiments are performed on two datasets: a private non-small cell lung cancer (NSCLC) dataset and AutoPET [132], a whole-body PET/CT dataset from MICCAI 2022 challenge.

#### 5.1.6. Ophthalmology

XAI is also employed in the segmentation of ophthalmological images. Optic disc and cup segmentation is explored in the setting of continual learning [136], where it is investigated in multi-site fundus datasets. Importance Activation Mapping is used to visualize the memorized content, facilitating an explanation of the model's memory. The focus is on reducing the model's forgetting. In [146], Seg-Grad-CAM is applied to ophthalmology for segmenting retinal layer boundaries. The study provides an entropy-based uncertainty visualization of segmentation probabilities. This offers more information about which retinal layers and regions exhibit higher uncertainty and allows for focusing on problematic areas. It is observed that higher uncertainty is associated with segmentation errors once it reaches a certain threshold. The experiments are performed on NR206.<sup>3</sup>

#### 5.1.7. Pancreatology

In [26], an interpretable image segmentation approach is proposed for pancreas segmentation in CT scans. The method is also compared to Grad-CAM and occlusion sensitivity, demonstrating its superior inference time. This method identifies regions in the input images where noise can be applied without significantly affecting model performance. It relies on noisy image occlusion and can be classified as a perturbation-based technique. To directly parameterize the noise mask for each pixel without harming the model's performance, an additional small interpretability model is trained. Both interpretability and utility models are based on U-Net. Pixels that can be significantly perturbed without changing the model's performance are considered less important. Essentially, the proposed method involves training noise distributions. This approach allows training dynamic noise maps for individual images, differing from the typical static systematic occlusion. Experiments are performed on a pancreas dataset [198]. In [143], a smoothing loss is introduced to guide interpretability learning. The authors observe that the explanations produced by U-Noise are less continuous. Assuming that important pixels are likely to be spatially close, the proposed smoothing objective considers the correlation between pixels during optimization. The resulting explanations are compared to those generated by Grad-CAM and U-Noise. Experiments are performed on a pancreas segmentation dataset [142] from the medical segmentation decathlon.

#### 5.1.8. Urology

In [154], a Bayesian approach is proposed to address the problem of interpreting domain-invariant features. The experiments are performed for prostate and cardiac segmentation tasks. The experiments are performed on T2 prostate

<sup>3</sup> <https://github.com/Medical-Image-Analysis/Retinal-layer-segmentation>.

MRI images from NCI-ISBI 2013 [147], I2CVB [148], and PROMISE12 [149]. For cardiac segmentation, MSCMR [150], EMIDEC [151], ACDC [102], MMWHS [152], and CASDC 2013 [153] datasets are used.

### 5.1.9. Anesthesiology

In [140], an interpretable approach is investigated for regional neuraxial analgesia monitoring. The experiments focus on thermal foot images for patients who have received epidural anesthesia. The proposed method is based on Convolutional Random Fourier Features (CRFF) and layer-wise weighted CAM. CRFF and CAM-based techniques are investigated in three segmentation models: U-Net, FCN, and Res-U-Net. CRRF gradient layers are added at skip connections. The initial results indicate that the integration of CRRF gradient layers allows to better differentiate background and foreground classes. The experiments are performed on the ThermalFeet<sup>4</sup> dataset of infrared images.

## 5.2. Industry-related applications

Various industrial and industry-related activities require precise segmentation. These activities might range from precise manufacturing and processing [90] to structural health monitoring in infrastructure, particularly in evaluating damage [199]. Industrial applications of XAI are also related to risk management. For example, the AI TriSM (Trust, Risk and Security Management) framework [200,201], has been adopted by several industrial organizations to enhance their risk management practices. This high-level framework includes both explainability and model monitoring as key components. Risk management is of particular importance in sensitive dynamic operational environments, where AI models need to be continuously assessed for performance and reliability.

In this section, we discuss both industrial processes and indirectly related tasks, such as environmental monitoring and remote sensing, which can have potential in industrial applications in a more narrow sense. Here, we divide industry-related explainable segmentation solutions into four categories: remote sensing, monitoring, scene understanding, and other more general applications.

### 5.2.1. Remote sensing

One of the first applications of interpretable image segmentation is in remote sensing. In [164], the U-Net model is applied for building detection. The proposed method works at the intersection of interpretability, representation learning, and interactive visualization, and is designed to explain U-Net's functionality. It employs Principal Component Analysis (PCA) on the activations in the bottleneck layer. PCA is the preferred method because it preserves the largest variance in the data. In the case of 3D visualizations, the first three components could be used. Following PCA, the new representations are clustered using k-means and DBSCAN algorithm. This approach allows for the visualization of learned

<sup>4</sup> <https://gcpds-image-segmentation.readthedocs.io/en/latest/notebooks/02-datasets.html>.

latent representations for all samples through an Intersection over Union (IoU)-based heatmap, allowing users to identify qualitatively different regions. The experiments are performed on Inria Aerial Image Labeling (IAIL) [163] dataset. This technique can be applied to detect and evaluate damages in industrial disasters or humanitarian crises, extending beyond mere infrastructure and product monitoring in industry. Another remote sensing application [178], specifically focusing on high-resolution satellite images, employs a gradient-free Sobol method [202] and a U-Noise model [26]. The proposed method is also compared to Seg-Grad-CAM++ classification extension.

### 5.2.2. Monitoring

Here we review relevant papers that offer explainable segmentation-based monitoring in proximate environments. In [90], simple gradient [24] saliency maps and SmoothGrad-based [43] saliencies are implemented for semantic segmentation models to investigate the adversarial attack setting. The experiments are performed on two industry-related cyber-physical system datasets. A private dataset from CTI FoodTech, a manufacturer of fruit-pitting machines, is used. In [82], the same private dataset is used for experiments with gradient-free XAI technique, based on the perturbations of intermediate activation maps.

In [176], the focus is on crack segmentation in critical infrastructures, such as tunnels and pavements. The U-Net model is used together with Grad-CAM, which is applied at the bottleneck, as in [16]. They investigate both simple and complex crack patterns as well as different backgrounds. Two other papers [199,203] also investigate the segmentation of different crack types. However, the proposed XAI techniques are implemented in classification models, and used for weakly supervised segmentation. These techniques are discussed in the subsequent section. In [175], an interpretable Bayesian network is used for facial micro-expression recognition. The authors prefer these networks for segmentation over DL models, primarily because of their superior causal interpretability when dealing with uncertain information. This can make them better interpretability candidates when uncertain causal inference is involved. The experiments are performed on the database [174] of face images.

The trade-off between the noisiness of the explanation and its computation time is particularly important in monitoring. Gradient-based methods are suitable for monitoring environments that require low latency and rapid decision-making. However, if the monitoring task allows for more time, slower non-gradient methods, such as perturbation-based explanations, might be a better choice.

### 5.2.3. Scene understanding

Scene understanding is an important area in applications for autonomous vehicles, monitoring of pedestrians and ambient objects, and surveillance. Precise real-time segmentation of road signs and obstructions is of particular importance. Explainable segmentation can be seen as part of explainable autonomous driving systems [167], which investigate events,

environments, and engine operations. An explainable variational autoencoder (VAE) model is proposed in [167], focusing on neuron activations with the use of attention mapping. For the experiments, the SYNTHIA [165] and A2D2 [166] datasets are used. The results are analyzed both qualitatively and quantitatively, using the average area under the receiver operator characteristic curve (AUC-ROC) index. In [171], XAI techniques are employed to investigate pixel-wise road detection for autonomous vehicles. The experiments are performed on different segmentation models, using the KITTI [170] road dataset. The problem is formulated as a binary segmentation task, where the classes are limited to the road and its surroundings. Grad-CAM and saliency maps are used to generate explanations. Unmanned aerial vehicles can also fall under the category of autonomous driving systems. In [90], gradient-based XAI techniques are applied to semantic drone dataset<sup>5</sup> from Graz University of Technology.

Automated semantic understanding of pedestrian environments is investigated in [169]. Here the focus is not on a particular XAI technique, but on introducing some level of explainability to segmentation evaluation. The paper argues that popular pixel-wise segmentation metrics, such as IoU or Dice coefficient, do not sufficiently take into account region-based over- and under-segmentation. Here over-segmentation refers to those cases where the relevant ground-truth region is segmented into a lower number of regions than the predicted mask. For instance, where there is only one bus in the segmented ground-truth, but the model segments it into three disjoint segments. In the case of under-segmentation, the opposite is true. Pixel-wise metrics do not accurately represent these differences in disjoint and joint regions as long as a large enough number of similar pixels is segmented in both the ground-truth image, and the corresponding prediction. The use of region-wise measures is proposed as a better way to explain the source of error in segmentation. The experiments are performed on PASCAL VOC 2012 [65], ADE20K [168], and Cityscapes [66]. In [204], the focus is on automatic semantic segmentation for sediment core analysis. To interpret the results, higher segmentation error regions and model prediction confidence are visualized. Here, the model confidence is defined as prediction probability, and the model error calculation is based on the normalized categorical cross-entropy.

The authors of [177] propose the Concept Relevance Propagation-based approach L-CRP as an extension of CRP [205]. By utilizing concept-based explanations, the study seeks to gain insights into both global and local aspects of explainability. The proposed approach seeks to understand the contribution of latent concepts to particular detections by identifying them, finding them in the input space, and evaluating their effect on relevance. Context scores are computed for different concepts. The experiments are performed on Cityscapes [66], Pascal VOC [65], and COCO [38] datasets.

#### 5.2.4. General applications

Some of the XAI-related experiments focus on more general datasets, typically used in evaluating the performance of segmentation models. COCO [38] dataset has been used as a benchmark in [82,89]. The dataset is composed of 21 classes of everyday objects, including several types of vehicles. Both [82,89] apply perturbation-based gradient-free methods. Input perturbations are used in [89], while feature map perturbations in pre-selected intermediate layers are used in [82].

Tendency-and-Assignment Explainer (TAX) framework is introduced in [125]. It seeks to explain: (1) what contributes to the segmentation output; and (2) what contributes to us thinking so (i.e. the why question). For this, a multi-annotator scenario is considered. The learned annotator-dependent prototype bank indicates the segmentation tendency, with a particular focus on uncertain regions. The experimental results on the Pascal VOC [65] dataset demonstrate that TAX predicts oversegmentation consistent with the annotator tendencies.

#### 5.3. XAI applications in self-supervised and weakly supervised segmentation

Manual image labeling is an expensive operation, especially when pixel-wise labeling is involved. It requires significant time and financial resources, and depending on the dataset being annotated, may also require particularly narrow expertise. With this in mind, it has been suggested that XAI techniques could be employed for automated labeling, which could also help reduce some forms of annotation bias.

In [206], a new explainable transformer architecture is proposed for model-inherent interpretability. The proposed model, a Siamese network, is investigated for weakly supervised segmentation. For enhanced interpretability, model representations are regularized using an attribute-guided loss function. Higher-layer attention maps are fused and used alongside attribute features. Qualitative segmentation results are compared with the SIPE technique. However, the model's limitation is its inability to incorporate attribute-level ground truth labels. Another application for weakly supervised segmentation employs LRP-based classification explanations [203]. These explanations are used to generate pixel-wise binary segmentations, which are then thresholded. The experiments are conducted on two datasets: one for cracks in sewer pipes and another for cracks in magnetic tiles [207].

In [199], surface crack detection and growth monitoring are investigated as part of structural health evaluation in infrastructure. Although no specific technique for explainable segmentation is proposed, explainable classification is used for weakly supervised segmentation, allowing for the quantification of crack severity. Six post-hoc techniques are implemented: InputXGradient, LRP, Integrated Gradients, DeepLift, DeepLift-Shap, and GradientShap. Additionally, B-cos networks and Neural Network Explainer are employed. In [208], GradCAM is used for semantic segmentation. An additional classification model is used with masked inputs, based on the given class. The classifier is trained on all the masked images

<sup>5</sup> <http://dronedataset.icg.tugraz.at/>.



across all classes. Explainable classification can also be used to enhance the data efficiency of segmentation models. For instance, in [209], Grad-CAM is employed to extract data-efficient features from the classification model, which are then used for segmentation. The results indicate that this approach generalizes across different segmentation methods.

## 6. Discussion

### 6.1. Open issues

Plenty of unresolved challenges remain in explainable semantic segmentation, most of which are also applicable to image classification tasks. Below is a non-exhaustive list of these challenges:

- **Evaluation metrics for XAI** Most of the literature on XAI in image classification focuses on introducing new explainability techniques and their modifications, rather than proposing new evaluative frameworks or benchmark datasets. This tendency is even more visible in explainable semantic segmentation. Currently, there are no papers dedicated solely to evaluating XAI results in image segmentation. The investigation of XAI metrics remains limited to the experimental results sections, and only in those few cases where quantitative evaluation is used. There is no consensus on which evaluation metrics are most crucial for capturing the key aspects of explainability, largely due to the difficulty in formalizing explainability-related concepts. A better theoretical understanding of the problem should inform the creation of evaluative XAI metrics and benchmarks. Such foundations would likely result in more efficient explainable segmentation methods that are better adapted to the problem at hand.
- **Safety and robustness of XAI methods** With the rapid deployment of DL models in medical, military, and industrial settings, XAI techniques are set to play an even more important role. Their primary use is driven by the need to determine if the model is reliable and trustworthy. However, similar questions can also be raised about the XAI techniques themselves. It is important to investigate their vulnerabilities and loopholes. Both deployers and end-users need to know whether they are secure against intentional attacks directed at XAI techniques or the model. Even if there is no direct threat, the robustness of each specific XAI method needs to be investigated on a case-by-case basis.  
Just like classification models, semantic segmentation models are susceptible to adversarial attacks. Different attack methods have been proposed [210–212]. When discussing adversarial attacks, it is common to focus on the model’s output as the primary target. However, it is also possible to attack the output’s explanation saliency while leaving both the input and the output perceptibly unchanged. Such attacks have been introduced and investigated in the context of image classification [213].

It has also been demonstrated that these second-level attacks can be extended to image segmentation [90]. More research is needed to find the best ways to combat them, especially since new adversarial attacks are constantly being developed, and comprehensive safety guarantees are challenging to ensure. Systematic investigations need to be undertaken for both white-box attacks, where the attacked model is known to the attacker, and black-box attacks, where it is unknown. Similar investigations into the robustness of interpretable segmentation could contribute to the overall security of AI systems.

Adversarial examples are typically not part of the training and testing datasets. This omission can lead to vulnerabilities in deployed models. Another critical issue is the presence of biases. When the most salient regions of the explanation map fall outside the boundaries of the object of interest, this might signal not just a misguided prediction but also the potential presence of adversarial influences [15]. Natural adversarial examples [214] and their influence on XAI in segmentation could be investigated as well.

- **XAI for video segmentation** As semantic scene segmentation is not limited to 2D images, new interpretability techniques could be investigated for video data, where temporal semantic segmentation is carried out. Video object segmentation requires significantly more computational resources. To our knowledge, there are currently no studies investigating explainable image segmentation in a dynamic setting. The nature of dynamic scenes could introduce novel challenges not previously encountered in 2D segmentation contexts. For instance, one would need to add an additional temporal explanation axis to account for differences in interpretability maps across video frames. This task could be further extended to real-time semantic segmentation by focusing on ways to reduce the latency of the generated explanations.
- **Computational complexity** Different XAI techniques require different computational resources, which may not be readily available in certain environments. Deployment constraints may include issues related to both software and hardware, where real-time services have to be ensured for edge devices and online service platforms [215]. Further experimental studies should investigate methods for reducing the computational complexity involved in generating explanations. This includes evaluating and optimizing the trade-offs between explanation quality and generation latency, particularly within diverse industrial contexts. Post-hoc methods, with perturbation-based techniques in particular, are rather inefficient in terms of explanation generation time. The total cost of generating local explanations increases with each new input image that requires interpretation, making it crucial to understand your use case and carefully assess available resources when selecting an XAI technique. Further research can help develop more efficient and scalable XAI methods, ensuring that high-quality explanations are provided within the resource constraints of specific applications.



## 6.2. Future directions

- **XAI benchmarking and evaluations** Given that most literature primarily focuses on qualitative metrics, we aim to highlight the need for a well-defined benchmark and evaluation strategy for XAI methods in image segmentation. To our knowledge, there are currently no studies on evaluations or benchmarks specifically for XAI methods in this area. Moreover, research focusing on the formal aspects of quantitative metrics in XAI is limited.
- **Mechanistic interpretability and other XAI approaches** We see mechanistic interpretability [216] as a promising research area. This approach seeks to reverse-engineer how models function. As far as we are aware, there have been no significant contributions in formal explainability [217] or argumentative XAI [218] within the context of image segmentation.
- **XAI for transformers** Additionally, there has not been much research [219] into the interpretability of transformers for segmentation, especially compared to convolutional networks. This is of particular interest given the growing popularity of transformer architectures in various applications. Conducting more comparative studies between XAI techniques for different architectures or different convolutional operations, such as atrous convolutions, could also be explored.
- **Failure Modes** This area is related to evaluation metrics. However, it covers problematic areas that could not be identified by the commonly used metrics. Specifically, XAI could be used to identify and mitigate bias in segmentation models. A systematic analysis of failure cases and potential failure modes could better determine the scope of applicability for XAI methods. Several studies [39] have critically evaluated different groups of explainability techniques in classification. However, a similar investigation has not yet been conducted in image segmentation.
- **Neural architecture search** Neural architecture search (NAS) explores automating neural architecture designs. XAI techniques can be applied in NAS in at least two distinct ways. First, existing XAI methods can be incorporated into NAS algorithms to improve their performance. For example, in [220], an explainable CAM technique is integrated with the NAS algorithm to avoid fully training submodels. Second, NAS algorithms can include interpretability aspects as one of the metrics to be optimized in multi-objective optimization. In [221], a surrogate interpretability metric has been used for multi-objective optimization in image classification. However, currently, no similar approaches exist for semantic segmentation tasks.
- **Continual Learning** Continual learning (CL) refers to the research area that investigates techniques allowing models to learn new tasks without forgetting the previously learned ones. This strong tendency for DL models

to forget previously learned information upon acquiring new knowledge is commonly described as catastrophic forgetting. More efficient solutions to CL problems would allow the models to be used more resourcefully, without retraining them from scratch when new data arrives. The intersection of XAI and CL presents an interesting area for investigation. XAI methods can be employed in CL to: (1) improve the model's performance; (2) better understand and explain the model's predictions; and (3) investigate the phenomenon of catastrophic forgetting. The exploration of XAI and CL could also lead to improved model understanding when either a shift in data distribution or concept drift occurs.

## 7. Conclusion

This survey presents a comprehensive view of the field of XAI in image segmentation. Our goal has been twofold: first, to provide an up-to-date literature review of various types of interpretability methods applied in semantic segmentation; and second, to clarify conceptual misunderstandings by proposing a method-centered taxonomy for image segmentation and general frameworks for different types of interpretability techniques. To these ends, we have categorized the methods into five major subgroups: prototype-based, gradient-based, perturbation-based, counterfactual methods, and architecture-based techniques. Based on the surveyed literature on explainable image segmentation, it is evident that most of the methods focus on local explanations and rely on qualitative evaluation. We hope this work can benefit computer vision researchers by presenting the landscape of XAI in image segmentation, delineating clearer boundaries between existing methods, and informing the development of new interpretability techniques.

## Acronyms

The following acronyms are used throughout this paper:

BUS	Breast ultrasound
CAM	Class activation mapping
CL	Continual learning
CMRI	Cardiac magnetic resonance imaging
CT	Computed tomography
IR	Infrared
LRP	Layer-wise relevance propagation
MG	Mammography
MRA	Magnetic resonance angiography
MRI	Magnetic resonance imaging
NAS	Neural architecture search
NIR	Near-infrared imaging
OCT	Optical coherence tomography
PET	Positron emission tomography
US	Ultrasound
WSI	Whole slide imaging

## CRedit authorship contribution statement

**Rokas Gipiškis:** Writing – review & editing, Writing – original draft, Visualization, Investigation, Conceptualization.

**Chun-Wei Tsai:** Writing – review & editing, Visualization, Supervision, Conceptualization. **Olga Kurasova:** Writing – review & editing, Supervision, Conceptualization.

### Declaration of competing interest

The authors declare that there is no conflict of interest in this paper.

### Acknowledgments

The authors would like to thank the editor and anonymous reviewers for their valuable comments and suggestions on the paper. This work was supported in part by the National Science and Technology Council of Taiwan, R.O.C., under Contract NSTC111-2222-E-110-006-MY3, NSTC112-2628-E-110-001-MY3, and NSTC113-2634-F-110-001-MBK.

### References

- [1] T. Miller, Explanation in artificial intelligence: Insights from the social sciences, *Artificial Intelligence* 267 (2019) 1–38.
- [2] C. Rudin, Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead, *Nat. Mach. Intell.* 1 (5) (2019) 206–215.
- [3] R. Roscher, B. Bohn, M.F. Duarte, J. Garcke, Explainable machine learning for scientific insights and discoveries, *IEEE Access* 8 (2020) 42200–42216.
- [4] E.H. Shortliffe, B.G. Buchanan, A model of inexact reasoning in medicine, *Math. Biosci.* 23 (3–4) (1975) 351–379.
- [5] W. Swartout, C. Paris, J. Moore, Explanations in knowledge systems: Design for explainable expert systems, *IEEE Expert* 6 (3) (1991) 58–64.
- [6] J.H. Friedman, Greedy function approximation: A gradient boosting machine, *Ann. Statist.* (2001) 1189–1232.
- [7] A. Krizhevsky, I. Sutskever, G.E. Hinton, ImageNet classification with deep convolutional neural networks, in: *Proceedings of the Advances in Neural Information Processing Systems*, Vol. 25, 2012, pp. 1–9.
- [8] D. Amodei, C. Olah, J. Steinhardt, P. Christiano, J. Schulman, D. Mané, Concrete problems in AI safety, 2016, pp. 1–29, arXiv preprint arXiv:1606.06565.
- [9] N. Maslej, L. Fattorini, E. Brynjolfsson, J. Etchemendy, K. Ligett, T. Lyons, J. Manyika, H. Ngo, J.C. Niebles, V. Parli, et al., Artificial intelligence index report 2023, 2023, pp. 1–386, arXiv preprint arXiv:2310.03715.
- [10] Morning Consult, IBM global AI adoption index - enterprise report, 2023, Available Online at: [https://filecache.mediaroom.com/mr5mr\\_ibmspgi/179414/download/IBM%20Global%20AI%20Adoption%20Index%20Report%20Dec.%202023.pdf](https://filecache.mediaroom.com/mr5mr_ibmspgi/179414/download/IBM%20Global%20AI%20Adoption%20Index%20Report%20Dec.%202023.pdf). (Accessed 24 April 2024).
- [11] C. Curtis, N. Gillespie, S. Lockey, AI-deploying organizations are key to addressing ‘perfect storm’ of AI risks, *AI Ethics* 3 (1) (2023) 145–153.
- [12] E. Commission, Regulation of the European parliament and of the council laying down harmonised rules on artificial intelligence (artificial intelligence act) and amending certain union legislative acts, 2021, Available Online at: <https://eur-lex.europa.eu/legal-content/EN/ALL/?uri=celex%3A52021PC0206>. (Accessed 22 April 2024).
- [13] The White House, Executive order on the safe, secure, and trustworthy development and use of artificial intelligence, 2023, Available Online at: <https://www.whitehouse.gov/briefing-room/presidential-actions/2023/10/30/executive-order-on-the-safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence/>. (Accessed 22 April 2024).
- [14] K. Wickstrøm, M. Kampffmeyer, R. Jenssen, Uncertainty and interpretability in convolutional neural networks for semantic segmentation of colorectal polyps, *Med. Image Anal.* 60 (2020) 101619.
- [15] L. Hoyer, M. Munoz, P. Katiyar, A. Khoreva, V. Fischer, Grid saliency for context explanations of semantic segmentation, in: *Proceedings of the Advances in Neural Information Processing Systems*, Vol. 32, 2019, pp. 1–12.
- [16] K. Vinogradova, A. Dibrov, G. Myers, Towards interpretable semantic segmentation via gradient-weighted class activation mapping (student abstract), in: *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34, 2020, pp. 13943–13944, no. 10.
- [17] D. Feng, C. Haase-Schütz, L. Rosenbaum, H. Hertlein, C. Glaeser, F. Timm, W. Wiesbeck, K. Dietmayer, Deep multi-modal object detection and semantic segmentation for autonomous driving: Datasets, methods, and challenges, *IEEE Trans. Intell. Transp. Syst.* 22 (3) (2020) 1341–1360.
- [18] S. Asgari Taghanaki, K. Abhishek, J.P. Cohen, J. Cohen-Adad, G. Hamarneh, Deep semantic segmentation of natural and medical images: A review, *Artif. Intell. Rev.* 54 (2021) 137–178.
- [19] K. Vinogradova, Explainable Artificial Intelligence for Image Segmentation and for Estimation of Optical Aberrations (Ph.D. dissertation), Dresden University of Technology, Germany, 2023, pp. 1–155.
- [20] S. Mullan, Deep Learning and Explainable AI in Medical Image Segmentation (Ph.D. dissertation), The University of Iowa, 2023, pp. 1–154.
- [21] S.N. Hasany, F. Mériaudeau, C. Petitjean, Post-hoc XAI in medical image segmentation: The journey thus far, in: *Proceedings of the Medical Imaging with Deep Learning*, 2024, pp. 1–17.
- [22] C. Molnar, *Interpretable Machine Learning: A Guide for Making Black Box Models Explainable*, Independently Published, 2022.
- [23] Y. Zhang, P. Tiño, A. Leonardi, K. Tang, A survey on neural network interpretability, *IEEE Trans. Emerg. Top. Comput. Intell.* 5 (5) (2021) 726–742.
- [24] K. Simonyan, A. Vedaldi, A. Zisserman, Deep inside convolutional networks: Visualising image classification models and saliency maps, 2013, pp. 1–8, arXiv preprint arXiv:1312.6034.
- [25] R.R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, D. Batra, Grad-cam: Visual explanations from deep networks via gradient-based localization, in: *Proceedings of the IEEE International Conference on Computer Vision*, 2017, pp. 618–626.
- [26] T. Koker, F. Mireshghallah, T. Titcombe, G. Kaissis, U-noise: Learnable noise masks for interpretable image segmentation, in: *2021 IEEE International Conference on Image Processing*, 2021, pp. 394–398.
- [27] M.D. Zeiler, R. Fergus, Visualizing and understanding convolutional networks, in: *Proceedings of the European Conference on Computer Vision*, 2014, pp. 818–833.
- [28] M.T. Ribeiro, S. Singh, C. Guestrin, “Why should I trust you?” explaining the predictions of any classifier, in: *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2016, pp. 1135–1144.
- [29] S.M. Lundberg, S.-I. Lee, A unified approach to interpreting model predictions, in: *Proceedings of the Advances in Neural Information Processing Systems*, Vol. 30, 2017, pp. 4768–4777.
- [30] V. Petsiuk, A. Das, K. Saenko, RISE: Randomized input sampling for explanation of black-box models, 2018, pp. 1–17, arXiv preprint arXiv:1806.07421.
- [31] D. Erhan, Y. Bengio, A. Courville, P. Vincent, Visualizing Higher-Layer Features of a Deep Network, *Tech. Rep.*, University of Montreal, 2009, p. 1341.
- [32] C. Olah, A. Mordvintsev, L. Schubert, Feature visualization, *Distill* 2 (11) (2017) e7.
- [33] O. Ronneberger, P. Fischer, T. Brox, U-net: Convolutional networks for biomedical image segmentation, in: *Proceedings of the 18th International Conference on Medical Image Computing and Computer-Assisted Intervention*, 2015, pp. 234–241.

- [34] V. Badrinarayanan, A. Kendall, R. Cipolla, SegNet: A deep convolutional encoder-decoder architecture for image segmentation, *IEEE Trans. Pattern Anal. Mach. Intell.* 39 (12) (2017) 2481–2495.
- [35] R. Strudel, R. Garcia, I. Laptev, C. Schmid, Segmenter: Transformer for semantic segmentation, in: *Proceedings of the IEEE/CVF International Conference on Computer Vision, 2021*, pp. 7262–7272.
- [36] J. Chen, Y. Lu, Q. Yu, X. Luo, E. Adeli, Y. Wang, L. Lu, A.L. Yuille, Y. Zhou, TransUNet: Transformers make strong encoders for medical image segmentation, 2021, pp. 1–13, arXiv preprint arXiv:2102.04306.
- [37] S.S. Kim, N. Meister, V.V. Ramaswamy, R. Fong, O. Russakovsky, Hive: evaluating the human interpretability of visual explanations, in: *Proceedings of the European Conference on Computer Vision, 2022*, pp. 280–298.
- [38] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, C.L. Zitnick, Microsoft COCO: Common objects in context, in: *Proceedings of the ECCV 2014, Zurich, Switzerland, September 6–12, Part V 13, 2014*, pp. 740–755.
- [39] J. Adebayo, J. Gilmer, M. Muelly, I. Goodfellow, M. Hardt, B. Kim, Sanity checks for saliency maps, in: *Proceedings of the Advances in Neural Information Processing Systems, Vol. 31, 2018*, pp. 1–11.
- [40] P.-J. Kindermans, S. Hooker, J. Adebayo, M. Alber, K.T. Schütt, S. Dähne, D. Erhan, B. Kim, The (un) reliability of saliency methods, in: *Explainable AI: Interpreting, Explaining and Visualizing Deep Learning, 2019*, pp. 267–280.
- [41] A. Saporta, X. Gui, A. Agrawal, A. Pareek, S.Q. Truong, C.D. Nguyen, V.-D. Ngo, J. Seekins, F.G. Blankenberg, A.Y. Ng, et al., Benchmarking saliency methods for chest X-ray interpretation, *Nat. Mach. Intell.* 4 (10) (2022) 867–878.
- [42] B. Bilodeau, N. Jaques, P.W. Koh, B. Kim, Impossibility theorems for feature attribution, 2022, pp. 1–38, arXiv preprint arXiv:2212.11870.
- [43] D. Smilkov, N. Thorat, B. Kim, F. Viégas, M. Wattenberg, Smoothgrad: removing noise by adding noise, 2017, pp. 1–10, arXiv preprint arXiv:1706.03825.
- [44] M. Ancona, E. Ceolini, C. Öztireli, M. Gross, Towards better understanding of gradient-based attribution methods for deep neural networks, 2017, pp. 1–16, arXiv preprint arXiv:1711.06104.
- [45] J. Adebayo, M. Muelly, H. Abelson, B. Kim, Post hoc explanations may be ineffective for detecting unknown spurious correlation, in: *Proceedings of the International Conference on Learning Representations, 2021*, pp. 1–13.
- [46] A. Atrey, K. Clary, D. Jensen, Exploratory not explanatory: Counterfactual analysis of saliency maps for deep reinforcement learning, 2019, pp. 1–23, arXiv preprint arXiv:1912.05743.
- [47] G. Nguyen, D. Kim, A. Nguyen, The effectiveness of feature attribution methods and its correlation with automatic evaluation scores, in: *Proceedings of the Advances in Neural Information Processing Systems, Vol. 34, 2021*, pp. 26422–26436.
- [48] Y. Zhou, S. Booth, M.T. Ribeiro, J. Shah, Do feature attribution methods correctly attribute features? in: *Proceedings of the AAAI Conference on Artificial Intelligence, Vol. 36, 2022*, pp. 9623–9633, no. 9.
- [49] R. Geirhos, R.S. Zimmermann, B. Bilodeau, W. Brendel, B. Kim, Don't trust your eyes: on the (un) reliability of feature visualizations, 2023, pp. 1–32, arXiv preprint arXiv:2306.04719.
- [50] A.B. Arrieta, N. Díaz-Rodríguez, J. Del Ser, A. Bannetot, S. Tabik, A. Barbado, S. García, S. Gil-López, D. Molina, R. Benjamins, R. Chatila, F. Herrera, Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI, *Inf. Fusion* 58 (2020) 82–115.
- [51] F.-L. Fan, J. Xiong, M. Li, G. Wang, On interpretability of artificial neural networks: A survey, *IEEE Trans. Radiat. Plasma Med. Sci.* 5 (6) (2021) 741–760.
- [52] M. Graziani, L. Dutkiewicz, D. Calvaresi, J.P. Amorim, K. Yordanova, M. Vered, R. Nair, P.H. Abreu, T. Blanke, V. Pulignano, J.O. Prior, L. Lauwaert, W. Reijers, A. Depeursinge, V. Andrearczyk, H. Müller, A global taxonomy of interpretable AI: Unifying the terminology for the technical and social sciences, *Artif. Intell. Rev.* 56 (4) (2023) 3473–3504.
- [53] T. Speith, A review of taxonomies of explainable artificial intelligence (XAI) methods, in: *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency, 2022*, pp. 2239–2250.
- [54] G. Schwalbe, B. Finzel, A comprehensive taxonomy for explainable artificial intelligence: A systematic survey of surveys on methods and concepts, *Data Min. Knowl. Discov.* (2023) 1–59.
- [55] M. Sacha, D. Rymarczyk, Ł. Struski, J. Tabor, B. Zieliński, ProtoSeg: Interpretable semantic segmentation with prototypical parts, in: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision, 2023*, pp. 1481–1492.
- [56] A. Shahroudnejad, A survey on understanding, visualizations, and explanation of deep neural networks, 2021, arXiv preprint arXiv:2102.01792.
- [57] M. Chromik, M. Schuessler, A taxonomy for human subject evaluation of black-box explanations in XAI, in: *ExSS-ATEC@IUI, Vol. 1, 2020*, pp. 1–7.
- [58] M. Biehl, B. Hammer, T. Villmann, Prototype-based models in machine learning, *Wiley Interdiscip. Rev.: Cogn. Sci.* 7 (2) (2016) 92–111.
- [59] C. Chen, O. Li, D. Tao, A. Barnett, C. Rudin, J.K. Su, This looks like that: Deep learning for interpretable image recognition, in: *Proceedings of the Advances in Neural Information Processing Systems, Vol. 32, 2019*, pp. 8930–8941.
- [60] J. Donnelly, A.J. Barnett, C. Chen, Deformable ProtoPNet: An interpretable image classifier using deformable prototypes, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2022*, pp. 10265–10275.
- [61] D. Rymarczyk, Ł. Struski, M. Górszczak, K. Lewandowska, J. Tabor, B. Zieliński, Interpretable image classification with differentiable prototypes assignment, in: *Proceedings of the European Conference on Computer Vision, 2022*, pp. 351–368.
- [62] Z. Zhang, P. Angelov, E. Soares, N. Longepe, P.P. Mathieu, An interpretable deep semantic segmentation method for earth observation, in: *Proceedings of the IEEE International Conference on Intelligent Systems, 2022*, pp. 1–8.
- [63] H. Jeffreys, *The Theory of Probability*, OUP Oxford, 1998.
- [64] J. Hilton, N. Cammarata, S. Carter, G. Goh, C. Olah, Understanding RL vision, *Distill* 5 (11) (2020) e29.
- [65] M. Everingham, L. Van Gool, C.K. Williams, J. Winn, A. Zisserman, The PASCAL visual object classes (VOC) challenge, *Int. J. Comput. Vis.* 88 (2010) 303–338.
- [66] M. Cordts, M. Omran, S. Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth, B. Schiele, The cityscapes dataset for semantic urban scene understanding, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2016*, pp. 3213–3223.
- [67] Segmentation of neuronal structures in EM stacks challenge- ISBI2012-imagej.net, 2012, <https://imagej.net/events/isbi-2012-segmentation-challenge>. (Accessed 4 July 2024).
- [68] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, A.L. Yuille, DeepLab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected CRFs, *IEEE Trans. Pattern Anal. Mach. Intell.* 40 (4) (2017) 834–848.
- [69] P. Angelov, E. Soares, Towards explainable deep neural networks (xDNN), *Neural Netw.* 130 (2020) 185–194.
- [70] D. Sculley, Web-scale k-means clustering, in: *Proceedings of the 19th International Conference on World Wide Web, 2010*, pp. 1177–1178.
- [71] G. Mateo-Garcia, J. Veitch-Michaelis, L. Smith, S.V. Oprea, G. Schumann, Y. Gal, A.G. Baydin, D. Backes, Towards global flood mapping onboard low cost satellites with machine learning, *Sci. Rep.* 11 (1) (2021) 7249.
- [72] S. Wachter, B. Mittelstadt, C. Russell, Counterfactual explanations without opening the black box: Automated decisions and the GDPR, *Harv. J. Law Technol.* 31 (2017) 841.



- [73] R. Guidotti, Counterfactual explanations and how to find them: Literature review and benchmarking, *Data Min. Knowl. Discov.* (2022) 1–55.
- [74] M. Zemni, M. Chen, É. Zablocki, H. Ben-Younes, P. Pérez, M. Cord, OCTET: Object-aware counterfactual explanations, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 15062–15071.
- [75] F. Yu, H. Chen, X. Wang, W. Xian, Y. Chen, F. Liu, V. Madhavan, T. Darrell, BDD100K: A diverse driving dataset for heterogeneous multitask learning, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 2636–2645.
- [76] Y. Xu, X. Yang, L. Gong, H.-C. Lin, T.-Y. Wu, Y. Li, N. Vasconcelos, Explainable object-induced action decision for autonomous vehicles, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 9523–9532.
- [77] M. Zemni, M. Chen, E. Zablocki, H. Ben-Younes, P. Pérez, M. Cord, OCTET: Object-aware counterfactual explanations—Supplementary material—, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 1–10.
- [78] D. Singh, A. Somani, A. Horsch, D.K. Prasad, Counterfactual explainable gastrointestinal and colonoscopy image segmentation, in: *Proceedings of the IEEE 19th International Symposium on Biomedical Imaging*, 2022, pp. 1–5.
- [79] D. Jha, P.H. Smedsrud, M.A. Riegler, P. Halvorsen, T. de Lange, D. Johansen, H.D. Johansen, Kvasir-seg: A segmented polyp dataset, in: *Proceedings of the MultiMedia Modeling Conference*, Daejeon, South Korea, January 5–8, Part II 26, 2020, pp. 451–462.
- [80] D. Jha, S. Ali, K. Emanuelsen, S.A. Hicks, V. Thambawita, E. Garcia-Ceja, M.A. Riegler, T. de Lange, P.T. Schmidt, H.D. Johansen, et al., Kvasir-instrument: Diagnostic and therapeutic tool segmentation dataset in gastrointestinal endoscopy, in: *Proceedings of the MultiMedia Modeling Conference*, Prague, Czech Republic, June 22–24, Part II 27, 2021, pp. 218–229.
- [81] P. Jacob, É. Zablocki, H. Ben-Younes, M. Chen, P. Pérez, M. Cord, STEEX: Steering counterfactual explanations with semantics, in: *Proceedings of the European Conference on Computer Vision*, 2022, pp. 387–403.
- [82] R. Gipiškis, D. Chiaro, D. Annunziata, F. Piccialli, Ablation studies in activation maps for explainable semantic segmentation in industry 4.0, in: *Proceedings of the IEEE EUROCON*, 2023, pp. 36–41.
- [83] S. Desai, H.G. Ramaswamy, Ablation-CAM: Visual explanations for deep convolutional network via gradient-free localization, in: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2020, pp. 983–991.
- [84] Y. LeCun, The MNIST database of handwritten digits, 1998, <http://yann.lecun.com/exdb/mnist/>.
- [85] A. Wan, D. Ho, Y. Song, H. Tillman, S.A. Bargal, J.E. Gonzalez, SegNBDT: Visual decision rules for segmentation, 2020, pp. 1–15, arXiv preprint arXiv:2006.06868.
- [86] R. Mottaghi, X. Chen, X. Liu, N.-G. Cho, S.-W. Lee, S. Fidler, R. Urtasun, A. Yuille, The role of context for object detection and semantic segmentation in the wild, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2014, pp. 891–898.
- [87] K. Gong, X. Liang, D. Zhang, X. Shen, L. Lin, Look into person: Self-supervised structure-sensitive learning and a new benchmark for human parsing, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 932–940.
- [88] P. Dardouillet, A. Benoit, E. Amri, P. Bolon, D. Dubucq, A. Crédoz, Explainability of image semantic segmentation through SHAP values, in: *Proceedings of the ICPR Workshops of the International Conference on Pattern Recognition Workshops*, 2022, pp. 188–202.
- [89] R. Gipiškis, O. Kurasova, Occlusion-based approach for interpretable semantic segmentation, in: *Proceedings of the Iberian Conference on Information Systems and Technologies*, CISTI, 2023, pp. 1–6.
- [90] R. Gipiškis, D. Chiaro, M. Preziosi, E. Prezioso, F. Piccialli, The impact of adversarial attacks on interpretable semantic segmentation in cyber-physical systems, *IEEE Syst. J.* (2023) 5327–5334.
- [91] J. Chiu, C.C. Li, O.J. Mengshoel, Potential applications of deep learning in automatic rock joint trace mapping in a rock mass, in: *Proceedings of the IOP Conference*, Vol. 1124, 2023, pp. 1–8, no. 1.
- [92] S.N. Hasany, C. Petitjean, F. Mériaudeau, Seg-XRes-CAM: Explaining Spatially Local Regions in image segmentation, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 3732–3737.
- [93] R.L. Draeos, L. Carin, Use hiresCAM instead of grad-CAM for faithful explanations of convolutional neural networks, 2020, pp. 1–20, arXiv preprint arXiv:2011.08891.
- [94] A.K. Gizzini, M. Shukor, A.J. Ghandour, Extending CAM-based XAI methods for remote sensing imagery segmentation, 2023, pp. 1–7, arXiv preprint arXiv:2310.01837.
- [95] S. Ji, S. Wei, M. Lu, Fully convolutional networks for multisource building extraction from an open aerial and satellite imagery data set, *IEEE Trans. Geosci. Remote Sens.* 57 (1) (2018) 574–586.
- [96] C. Schorr, P. Goodarzi, F. Chen, T. Dahmen, Neuroscope: An explainable ai toolbox for semantic segmentation and image classification of convolutional neural nets, *Appl. Sci.* 11 (5) (2021) 2199.
- [97] A. Kapischnikov, T. Bolukbasi, F. Viégas, M. Terry, XRAI: Better attributions through regions, in: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 4948–4957.
- [98] D. Vázquez, J. Bernal, F.J. Sánchez, G. Fernández-Esparrach, A.M. López, A. Romero, M. Drozdzal, A. Courville, et al., A benchmark for endoluminal scene segmentation of colonoscopy images, *J. Healthc. Eng.* (2017) 1–9.
- [99] P. Bilic, P. Christ, H.B. Li, E. Vorontsov, A. Ben-Cohen, G. Kaissis, A. Szeskin, C. Jacobs, G.E.H. Mamani, G. Chartrand, et al., The liver tumor segmentation benchmark (LiTS), *Med. Image Anal.* 84 (2023) 102680.
- [100] V. Couteaux, O. Nempont, G. Pizaine, I. Bloch, Towards interpretability of segmentation networks by analyzing DeepDreams, in: *Proceedings of the Second International Workshop, IMIMIC 2019, and 9th International Workshop, ML-CDS 2019, Held in Conjunction with MICCAI Conference*, Shenzhen, China, 2019, pp. 56–63.
- [101] P. Radau, Y. Lu, K. Connelly, G. Paul, A.J. Dick, G.A. Wright, Evaluation framework for algorithms segmenting short axis cardiac MRI, *MIDAS J.* (2009).
- [102] O. Bernard, A. Lalande, C. Zotti, F. Cervenansky, X. Yang, P.-A. Heng, I. Cetin, K. Lekadir, O. Camara, M.A.G. Ballester, et al., Deep learning techniques for automatic MRI cardiac multi-structures segmentation and diagnosis: Is the problem solved? *IEEE Trans. Med. Imaging* 37 (11) (2018) 2514–2525.
- [103] J. Sun, F. Darbehani, M. Zaidi, B. Wang, SAUNet: Shape attentive U-net for interpretable medical image segmentation, in: *Proceedings of the MICCAI Conference*, Lima, Peru, Part IV 23, 2020, pp. 797–806.
- [104] A. Santamaria-Pang, J. Kubricht, A. Chowdhury, C. Bhushan, P. Tu, Towards emergent language symbolic semantic segmentation and model interpretability, in: *Proceedings of the MICCAI Conference*, Lima, Peru, October 4–8, Part I 23, 2020, pp. 326–334.
- [105] N. Codella, V. Rotemberg, P. Tschandl, M.E. Celebi, S. Dusza, D. Gutman, B. Helba, A. Kalloo, K. Liopyris, M. Marchetti, et al., Skin lesion analysis toward melanoma detection 2018: A challenge hosted by the international skin imaging collaboration (ISIC), 2019, pp. 1–12, arXiv preprint arXiv:1902.03368.
- [106] R. Gu, G. Wang, T. Song, R. Huang, M. Aertsen, J. Deprest, S. Ourselin, T. Vercauteren, S. Zhang, CA-Net: Comprehensive attention convolutional neural networks for explainable medical image segmentation, *IEEE Trans. Med. Imaging* 40 (2) (2020) 699–711.
- [107] A. Janik, J. Dodd, G. Ifrim, K. Sankaran, K. Curran, Interpretability of a deep learning model in the application of cardiac MRI segmentation with an ACDC challenge dataset, in: *Proceedings of SPIE Medical Imaging Conference*, Vol. 11596, 2021, pp. 861–872.
- [108] A. Ahmed, L.A. Ali, Explainable medical image segmentation via generative adversarial networks and layer-wise relevance propagation, 2021, pp. 1–3, arXiv preprint arXiv:2111.01665.



- [109] B.H. Menze, A. Jakab, S. Bauer, J. Kalpathy-Cramer, K. Farahani, J. Kirby, Y. Burren, N. Porz, J. Slotboom, R. Wiest, et al., The multimodal brain tumor image segmentation benchmark (BRATS), *IEEE Trans. Med. Imaging* 34 (10) (2014) 1993–2024.
- [110] H. Saleem, A.R. Shahid, B. Raza, Visual interpretability in 3D brain tumor segmentation network, *Comput. Biol. Med.* 133 (2021) 104410.
- [111] M. Trokielewicz, A. Czajka, P. Maciejewicz, Post-mortem iris recognition resistant to biological eye decay processes, in: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2020, pp. 2307–2315.
- [112] A. Kuehlkamp, A. Boyd, A. Czajka, K. Bowyer, P. Flynn, D. Chute, E. Benjamin, Interpretable deep learning-based forensic iris segmentation and recognition, in: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2022, pp. 359–368.
- [113] P. Tschandl, C. Rosendahl, H. Kittler, The HAM10000 dataset, a large collection of multi-source dermatoscopic images of common pigmented skin lesions, *Sci. Data* 5 (1) (2018) 1–9.
- [114] A.E. Kavur, N.S. Gezer, M. Barış, S. Aslan, P.-H. Conze, V. Groza, D.D. Pham, S. Chatterjee, P. Ernst, S. Özkan, et al., CHAOS challenge-combined (CT-MR) healthy abdominal organ segmentation, *Med. Image Anal.* 69 (2021) 101950.
- [115] M. Karri, C.S.R. Annavarapu, U.R. Acharya, Explainable multi-module semantic guided attention based network for medical image segmentation, *Comput. Biol. Med.* 151 (2022) 106231.
- [116] F. Ambellan, A. Tack, M. Ehlke, S. Zachow, Automated segmentation of knee bone and cartilage combining statistical shape knowledge and convolutional neural networks: Data from the Osteoarthritis Initiative, *Med. Image Anal.* 52 (2019) 109–118.
- [117] K. Schulze, F. Peppert, C. Schütte, V. Sunkara, Chimeric U-net—modifying the standard U-net towards explainability, 2022, pp. 1–12, bioRxiv.
- [118] R.A. Zeineldin, M.E. Karar, Z. Elshaer, J. Coburger, C.R. Wirtz, O. Burgert, F. Mathis-Ullrich, Explainability of deep neural networks for MRI analysis of brain tumors, *Int. J. Comput. Assist. Radiol. Surg.* 17 (9) (2022) 1673–1683.
- [119] S. Chatterjee, A. Das, C. Mandal, B. Mukhopadhyay, M. Vipinraj, A. Shukla, R. Nagaraja Rao, C. Sarasaen, O. Speck, A. Nürnberger, TorchEsegeta: Framework for interpretability and explainability of image-based deep learning models, *Appl. Sci.* 12 (4) (2022) 1834.
- [120] F. Bardozzo, T. Collins, A. Forgione, A. Hostettler, R. Tagliaferri, StaSiS-net: A stacked and siamese disparity estimation network for depth reconstruction in modern 3D laparoscopy, *Med. Image Anal.* 77 (2022) 102380.
- [121] F. Bardozzo, M.D. Priscoli, T. Collins, A. Forgione, A. Hostettler, R. Tagliaferri, Cross X-AI: Explainable semantic segmentation of laparoscopic images in relation to depth estimation, in: *Proceedings of the IEEE International Joint Conference on Neural Networks*, 2022, pp. 1–8.
- [122] W. Al-Dhabyani, M. Goma, H. Khaled, A. Fahmy, Dataset of breast ultrasound images, *Data Brief* 28 (2020) 104863.
- [123] J. Wang, Y. Zheng, J. Ma, X. Li, C. Wang, J. Gee, H. Wang, W. Huang, Information bottleneck-based interpretable multitask network for breast cancer classification and segmentation, *Med. Image Anal.* 83 (2023) 102687.
- [124] S.G. Armato III, G. McLennan, L. Bidaut, M.F. McNitt-Gray, C.R. Meyer, A.P. Reeves, B. Zhao, D.R. Aberle, C.I. Henschke, E.A. Hoffman, et al., The lung image database consortium (LIDC) and image database resource initiative (IDRI): A completed reference database of lung nodules on CT scans, *Med. Phys.* 38 (2) (2011) 915–931.
- [125] Y.-C. Cheng, Z.-Y. Shiau, F.-E. Yang, Y.-C.F. Wang, TAX: Tendency-and-assignment explainer for semantic segmentation with multi-annotators, 2023, pp. 1–10, arXiv preprint arXiv:2302.09561.
- [126] J.A. Dowling, J. Sun, P. Pichler, D. Rivest-Hénault, S. Ghose, H. Richardson, C. Wratten, J. Martin, J. Arm, L. Best, et al., Automatic substitute computed tomography generation and contouring for magnetic resonance imaging (MRI)-alone external beam radiation therapy from standard MRI sequences, *Int. J. Radiat. Oncol. Biol. Phys.* 93 (5) (2015) 1144–1153.
- [127] W. Dai, S. Liu, C.B. Engstrom, S.S. Chandra, Explainable semantic medical image segmentation with style, 2023, pp. 1–12, arXiv preprint arXiv:2303.05696.
- [128] B. Landman, Z. Xu, J. Igelsias, M. Styner, T. Langerak, A. Klein, MICCAI multi-atlas labeling beyond the cranial vault—workshop and challenge, in: *Proceedings of the MICCAI Multi-Atlas Labeling beyond Cranial Vault—Workshop Challenge*, Vol. 5, 2015, p. 12.
- [129] M. Xian, Y. Zhang, H.-D. Cheng, F. Xu, K. Huang, B. Zhang, J. Ding, C. Ning, Y. Wang, A Benchmark for Breast Ultrasound Image Segmentation, *BUSIS, Infinite Study*, 2018.
- [130] T. Geertsma, *Ultrasoundcases*. info, 2014.
- [131] M. Karimzadeh, A. Vakanski, M. Xian, B. Zhang, Post-hoc explainability of BI-RADS descriptors in a multi-task framework for breast cancer detection and segmentation, 2023, pp. 1–11, arXiv preprint arXiv:2308.14213.
- [132] S. Gatidis, T. Hepp, M. Früh, C. La Fougère, K. Nikolaou, C. Pfannenber, B. Schölkopf, T. Küstner, C. Cyran, D. Rubin, A whole-body FDG-PET/CT dataset with manually annotated tumor lesions, *Sci. Data* 9 (1) (2022) 601.
- [133] S. Kang, Z. Chen, L. Li, W. Lu, X.S. Qi, S. Tan, Learning feature fusion via an interpretation method for tumor segmentation on PET/CT, *Appl. Soft Comput.* 148 (2023) 110825.
- [134] K. Wang, S. Yin, Y. Wang, S. Li, Explainable deep learning for medical image segmentation with learnable class activation mapping, in: *Proceedings of the 2023 2nd Asia Conference on Algorithms, Computing and Machine Learning*, 2023, pp. 210–215.
- [135] R. Sun, M. Rostami, Explainable artificial intelligence architecture for melanoma diagnosis using indicator localization and self-supervised learning, 2023, pp. 1–16, arXiv preprint arXiv:2303.14615.
- [136] J. Zhang, R. Gu, P. Xue, M. Liu, H. Zheng, Y. Zheng, L. Ma, G. Wang, L. Gu, S3R: Shape and semantics-based selective regularization for explainable continual segmentation across multiple sites, *IEEE Trans. Med. Imaging* (2023) 1–13.
- [137] I.C. Moreira, I. Amaral, I. Domingues, A. Cardoso, M.J. Cardoso, J.S. Cardoso, INbreast: Toward a full-field digital mammographic database, *Academic Radiol.* 19 (2) (2012) 236–248.
- [138] A. Farrag, G. Gad, Z.M. Fadlullah, M.M. Fouda, M. Alsabaan, An explainable AI system for medical image segmentation with preserved local resolution: Mammogram tumor segmentation, *IEEE Access* (2023) 125543–125561.
- [139] M. Dörrich, M. Hecht, R. Fietkau, A. Hartmann, H. Iro, A.-O. Gostian, M. Eckstein, A.M. Kist, Explainable convolutional neural networks for assessing head and neck cancer histopathology, *Diagn. Pathol.* 18 (1) (2023) 121.
- [140] J.C. Aguirre-Arango, A.M. Álvarez-Meza, G. Castellanos-Dominguez, Feet segmentation for regional analgesia monitoring using convolutional RFF and layer-wise weighted CAM interpretability, *Computation* 11 (6) (2023) 113.
- [141] S. He, Y. Feng, P.E. Grant, Y. Ou, Segmentation ability map: Interpret deep features for medical image segmentation, *Med. Image Anal.* 84 (2023) 102726.
- [142] M. Antonelli, A. Reinke, S. Bakas, K. Farahani, A. Kopp-Schneider, B.A. Landman, G. Litjens, B. Menze, O. Ronneberger, R.M. Summers, et al., The medical segmentation decathlon, *Nat. Commun.* 13 (1) (2022) 4128.
- [143] T. Okamoto, C. Gu, J. Yu, C. Zhang, Generating smooth interpretability map for explainable image segmentation, in: *Proceedings of the IEEE Global Conference on Consumer Electronics*, 2023, pp. 1023–1025.
- [144] J. Li, P. Jin, J. Zhu, H. Zou, X. Xu, M. Tang, M. Zhou, Y. Gan, J. He, Y. Ling, et al., Multi-scale GCN-assisted two-stage network for joint segmentation of retinal layers and discs in peripapillary OCT images, *Biomed. Opt. Express* 12 (4) (2021) 2204–2220.
- [145] S.J. Chiu, M.J. Allingham, P.S. Mettu, S.W. Cousins, J.A. Izatt, S. Farsiu, Kernel regression based segmentation of optical coherence tomography images with diabetic macular edema, *Biomed. Opt. Express* 6 (4) (2015) 1172–1194.

- [146] X. He, Y. Wang, F. Poiesi, W. Song, Q. Xu, Z. Feng, Y. Wan, Exploiting multi-granularity visual features for retinal layer segmentation in human eyes, *Front. Bioeng. Biotechnol.* 11 (2023) 1–14.
- [147] N. Bloch, A. Madabhushi, H. Huisman, J. Freymann, J. Kirby, M. Grauer, A. Enquobahrie, C. Jaffe, L. Clarke, K. Farahani, NCI-ISBI 2013 challenge: Automated segmentation of prostate structures, *Cancer Imaging Arch.* 6 (2015) 5.
- [148] G. Lemaître, R. Martí, J. Freixenet, J.C. Vilanova, P.M. Walker, F. Meriaudeau, Computer-aided detection and diagnosis for prostate cancer based on mono and multi-parametric MRI: A review, *Comput. Biol. Med.* 60 (2015) 8–31.
- [149] G. Litjens, R. Toth, W. Van De Ven, C. Hoeks, S. Kerkstra, B. Van Ginneken, G. Vincent, G. Guillard, N. Birbeck, J. Zhang, et al., Evaluation of prostate segmentation algorithms for MRI: The PROMISE12 challenge, *Med. Image Anal.* 18 (2) (2014) 359–373.
- [150] X. Zhuang, Multivariate mixture model for myocardial segmentation combining multi-source images, *IEEE Trans. Pattern Anal. Mach. Intell.* 41 (12) (2018) 2933–2946.
- [151] A. Lalande, Z. Chen, T. Decourselle, A. Qayyum, T. Pommier, L. Lorgis, E. de La Rosa, A. Cochet, Y. Cottin, D. Ginjac, et al., Emidec: A database usable for the automatic evaluation of myocardial infarction from delayed-enhancement cardiac MRI, *Data* 5 (4) (2020) 89.
- [152] X. Zhuang, J. Shen, Multi-scale patch and multi-modality atlases for whole heart segmentation of MRI, *Med. Image Anal.* 31 (2016) 77–87.
- [153] H. Kirişli, M. Schaap, C. Metz, A. Dharampal, W.B. Meijboom, S.-L. Papadopoulou, A. Dedic, K. Nieman, M.A. de Graaf, M. Meijjs, et al., Standardized evaluation framework for evaluating coronary artery stenosis detection, stenosis quantification and lumen segmentation algorithms in computed tomography angiography, *Med. Image Anal.* 17 (8) (2013) 859–876.
- [154] S. Gao, H. Zhou, Y. Gao, X. Zhuang, BayeSeg: Bayesian modeling for medical image segmentation with interpretable generalizability, 2023, pp. 1–15, arXiv preprint arXiv:2303.01710.
- [155] X. He, W. Song, Y. Wang, F. Poiesi, J. Yi, M. Desai, Q. Xu, K. Yang, Y. Wan, Light-weight retinal layer segmentation with global reasoning, 2024, arXiv preprint arXiv:2404.16346.
- [156] Z. Lambert, C. Petitjean, B. Dubray, S. Kuan, SegTHOR: Segmentation of thoracic organs at risk in CT images, in: *Proceedings of the International Conference on Image Processing Theory, Tools and Applications*, 2020, pp. 1–6.
- [157] Z. Lambert, C. Le Guyader, About the incorporation of topological prescriptions in CNNs for medical image semantic segmentation, *J. Math. Imaging Vision* (2024) 1–28.
- [158] J. Sivaswamy, S. Krishnadas, A. Chakravarty, G. Joshi, A.S. Tabish, et al., A comprehensive retinal image dataset for the assessment of glaucoma from the optic nerve head analysis, *JSM Biomed. Imaging Data Pap.* 2 (1) (2015) 1004.
- [159] F. Fumero, S. Alayón, J.L. Sanchez, J. Sigut, M. Gonzalez-Hernandez, RIM-ONE: An open retinal image database for optic nerve evaluation, in: *Proceedings of the International Symposium on Computer-Based Medical Systems*, 2011, pp. 1–6.
- [160] J.I. Orlando, H. Fu, J.B. Breda, K. Van Keer, D.R. Bathula, A. Diaz-Pinto, R. Fang, P.-A. Heng, J. Kim, J. Lee, et al., Refuge challenge: A unified framework for evaluating automated methods for glaucoma assessment from fundus photographs, *Med. Image Anal.* 59 (2020) 101570.
- [161] M. Losch, M. Fritz, B. Schiele, Semantic bottlenecks: Quantifying and improving inspectability of deep representations, *Int. J. Comput. Vis.* 129 (2021) 3136–3153.
- [162] D. Bau, B. Zhou, A. Khosla, A. Oliva, A. Torralba, Network dissection: Quantifying interpretability of deep visual representations, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 6541–6549.
- [163] E. Maggiori, Y. Tarabalka, G. Charpiat, P. Alliez, Can semantic labeling methods generalize to any city? The inria aerial image labeling benchmark, in: *Proceedings of the IEEE International Geoscience and Remote Sensing Symposium*, 2017, pp. 3226–3229.
- [164] A. Janik, K. Sankaran, A. Ortiz, Interpreting black-box semantic segmentation models in remote sensing applications, *Mach. Learn. Methods Vis. Big Data* (2019) 7–11.
- [165] J. Zolfaghari Bengar, A. Gonzalez-Garcia, G. Villalonga, B. Raducanu, H. Habibi Aghdam, M. Mozerov, A.M. Lopez, J. Van de Weijer, Temporal coherence for active learning in videos, in: *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, 2019, pp. 1–10.
- [166] J. Geyer, Y. Kassahun, M. Mahmudi, X. Ricou, R. Durgesh, A.S. Chung, L. Hauswald, V.H. Pham, M. Mühlegg, S. Dorn, et al., A2d2: Audi autonomous driving dataset, 2020, pp. 1–10, arXiv preprint arXiv:2004.06320.
- [167] M. Abukmeil, A. Genovese, V. Piuri, F. Rundo, F. Scotti, Towards explainable semantic segmentation for autonomous driving systems by multi-scale variational attention, in: *Proceedings of the IEEE International Conference on Autonomous Systems*, 2021, pp. 1–5.
- [168] B. Zhou, H. Zhao, X. Puig, T. Xiao, S. Fidler, A. Barriuso, A. Torralba, Semantic understanding of scenes through the ADE20K dataset, *Int. J. Comput. Vis.* 127 (2019) 302–321.
- [169] Y. Zhang, S. Mehta, A. Caspi, Rethinking semantic segmentation evaluation for explainability and model selection, 2021, pp. 1–14, arXiv preprint arXiv:2101.08418.
- [170] J. Fritsch, T. Kuehnl, A. Geiger, A new performance measure and evaluation benchmark for road detection algorithms, in: *Proceedings of the IEEE Conference on Intelligent Transportation Systems*, 2013, pp. 1693–1700.
- [171] H. Mankodiya, D. Jadav, R. Gupta, S. Tanwar, W.-C. Hong, R. Sharma, OD-XAI: Explainable AI-based semantic object detection for autonomous vehicles, *Appl. Sci.* 12 (11) (2022) 5310.
- [172] C.-H. Lee, Z. Liu, L. Wu, P. Luo, MaskGAN: Towards diverse and interactive facial image manipulation, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 5549–5558.
- [173] Z. Liu, P. Luo, X. Wang, X. Tang, Deep learning face attributes in the wild, in: *Proceedings of the IEEE International Conference on Computer Vision*, 2015, pp. 3730–3738.
- [174] M. Turk, A. Pentland, Eigenfaces for recognition, *J. Cogn. Neurosci.* 3 (1) (1991) 71–86.
- [175] C. Wang, X. Gao, X. Li, An interpretable deep Bayesian model for facial micro-expression recognition, in: *Proceedings of the IEEE International Conference on Control and Robotics Engineering*, 2023, pp. 91–94.
- [176] F. Liu, W. Ding, Y. Qiao, L. Wang, Transfer learning-based encoder-decoder model with visual explanations for infrastructure crack segmentation: New open database and comprehensive evaluation, *Undergr. Space* (2023) 60–81.
- [177] M. Dreyer, R. Achibat, T. Wiegand, W. Samek, S. Lapuschkin, Revealing hidden context bias in segmentation and object detection through concept-specific explanations, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 3828–3838.
- [178] H. Shreim, A.K. Gizzini, A.J. Ghandour, Trainable noise model as an XAI evaluation method: application on sobol for remote sensing image segmentation, 2023, pp. 1–7, arXiv preprint arXiv:2310.01828.
- [179] R.R. Hoffman, S.T. Mueller, G. Klein, J. Litman, Metrics for explainable AI: Challenges and prospects, 2018, pp. 1–50, arXiv preprint arXiv:1812.04608.
- [180] M. Graziani, I. Palatnik de Sousa, M.M. Vellasco, E. Costa da Silva, H. Müller, V. Andrearczyk, Sharpening local interpretable model-agnostic explanations for histopathology: Improved understandability and reliability, in: *Proceedings of the MICCAI Conference*, Strasbourg, France, Part III 24, 2021, pp. 540–549.
- [181] C.-K. Yeh, C.-Y. Hsieh, A. Suggala, D.I. Inouye, P.K. Ravikumar, On the (in)fidelity and sensitivity of explanations, in: *Proceedings of the Advances in Neural Information Processing Systems*, 32, 2019, pp. 1–12.

- [182] N. Kohlikeyan, V. Miglani, M. Martin, E. Wang, B. Alsallakh, J. Reynolds, A. Melnikov, N. Kliushkina, C. Araya, S. Yan, et al., Captum: A unified and generic model interpretability library for pytorch, 2020, pp. 1–11, arXiv preprint [arXiv:2009.07896](https://arxiv.org/abs/2009.07896).
- [183] J. Colin, T. Fel, R. Cadène, T. Serre, What I cannot predict, I do not understand: A human-centered evaluation framework for explainability methods, in: *Proceedings of the Advances in Neural Information Processing Systems*, Vol. 35, 2022, pp. 2832–2845.
- [184] R. Tomsett, D. Harborne, S. Chakraborty, P. Gurram, A. Preece, Sanity checks for saliency metrics, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34, 2020, pp. 6021–6029, no. 04.
- [185] N.C. Codella, D. Gutman, M.E. Celebi, B. Helba, M.A. Marchetti, S.W. Dusza, A. Kalloo, K. Liopyris, N. Mishra, H. Kittler, et al., Skin lesion analysis toward melanoma detection: A challenge at the 2017 international symposium on biomedical imaging (ISBI), hosted by the international skin imaging collaboration (ISIC), in: *Proceedings of the IEEE International Symposium on Biomedical Imaging*, 2018, pp. 168–172.
- [186] M. Combalia, N.C. Codella, V. Rotemberg, B. Helba, V. Vilaplana, O. Reiter, C. Carrera, A. Barreiro, A.C. Halpern, S. Puig, et al., BCN20000: Dermoscopic lesions in the wild, 2019, pp. 1–3, arXiv preprint [arXiv:1908.02288](https://arxiv.org/abs/1908.02288).
- [187] B. Zhou, A. Khosla, A. Lapedriza, A. Oliva, A. Torralba, Learning deep features for discriminative localization, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 2921–2929.
- [188] J.T. Springenberg, A. Dosovitskiy, T. Brox, M. Riedmiller, Striving for simplicity: The all convolutional net, 2014, pp. 1–14, arXiv preprint [arXiv:1412.6806](https://arxiv.org/abs/1412.6806).
- [189] S. Mohagheghi, A.H. Foruzan, Developing an explainable deep learning boundary correction method by incorporating cascaded x-dim models to improve segmentation defects in liver CT images, *Comput. Biol. Med.* 140 (2022) 105106.
- [190] I. Dirks, M. Keyaerts, B. Neyns, J. Vandemeulebroucke, Computer-aided detection and segmentation of malignant melanoma lesions on whole-body 18F-FDG PET/CT using an interpretable deep learning approach, *Comput. Methods Programs Biomed.* 221 (2022) 106902.
- [191] S. Dasanayaka, S. Silva, V. Shantha, D. Meedeniya, T. Ambegoda, Interpretable machine learning for brain tumor analysis using MRI, in: *Proceedings of the IEEE International Conference on Advanced Research in Computing*, 2022, pp. 212–217.
- [192] K. Cortacero, B. McKenzie, S. Müller, R. Khazen, F. Lafouresse, G. Corsaut, N. Van Acker, F.-X. Frenois, L. Lamant, N. Meyer, et al., Evolutionary design of explainable algorithms for biomedical image segmentation, *Nature Commun.* 14 (1) (2023) 7112.
- [193] A. Kaur, G. Dong, A. Basu, GradXcepUNet: Explainable AI based medical image segmentation, in: *Proceedings of the International Conference on Smart Multimedia*, 2022, pp. 174–188.
- [194] P.F. Christ, F. Ettliger, F. Grün, M.E.A. Elshaera, J. Lipkova, S. Schlecht, F. Ahmaddy, S. Tatavarty, M. Bickel, P. Bilic, et al., Automatic liver and tumor segmentation of CT and MRI volumes using cascaded fully convolutional neural networks, 2017, pp. 1–20, arXiv preprint [arXiv:1702.05970](https://arxiv.org/abs/1702.05970).
- [195] E. Pintelas, I.E. Livieris, XSC—An explainable image segmentation and classification framework: A case study on skin cancer, *Electronics* 12 (17) (2023) 3551.
- [196] P. Ashtari, D.M. Sima, L. De Lathauwer, D. Sappéy-Marinié, F. Maes, S. Van Huffel, Factorizer: A scalable interpretable approach to context modeling for medical image segmentation, *Med. Image Anal.* 84 (2023) 102706.
- [197] H. Mattern, A. Sciarra, F. Godenschweger, D. Stucht, F. Lüsebrink, G. Rose, O. Speck, Prospective motion correction enables highest resolution time-of-flight angiography at 7T, *Magn. Reson. Med.* 80 (1) (2018) 248–258.
- [198] A.L. Simpson, M. Antonelli, S. Bakas, M. Bilello, K. Farahani, B. Van Ginneken, A. Kopp-Schneider, B.A. Landman, G. Litjens, B. Menze, et al., A large annotated medical image dataset for the development and evaluation of segmentation algorithms, 2019, pp. 1–15, arXiv preprint [arXiv:1902.09063](https://arxiv.org/abs/1902.09063).
- [199] F. Forest, H. Porta, D. Tuia, O. Fink, From classification to segmentation with explainable AI: A study on crack detection and growth monitoring, 2023, pp. 1–43, arXiv preprint [arXiv:2309.11267](https://arxiv.org/abs/2309.11267).
- [200] L. Goasduff, Gartner says CISOs need to champion AI TRISM to improve AI results, 2023, Available Online at: <https://www.gartner.com/en/newsroom/press-releases/2023-09-27-gartner-says-cisos-need-to-champion-ai-trism-to-improve-ai-results>. (Accessed 28 August 2024).
- [201] G. Aceto, D. Ciunzo, A. Montieri, V. Persico, A. Pescape, AI-powered internet traffic classification: Past, present, and future, *IEEE Commun. Mag.* (In Press, 2024).
- [202] T. Fel, R. Cadène, M. Chalvidal, M. Cord, D. Vigouroux, T. Serre, Look at the variance! efficient black-box explanations with Sobol-based sensitivity analysis, in: *Proceedings of the Advances in Neural Information Processing Systems*, Vol. 34, 2021, pp. 26005–26014.
- [203] C. Seibold, J. Künzel, A. Hilsmann, P. Eisert, From explanations to segmentation: Using explainable AI for image segmentation, 2022, pp. 1–10, arXiv preprint [arXiv:2202.00315](https://arxiv.org/abs/2202.00315).
- [204] A. Di Martino, G. Carlini, G. Castellani, D. Remondini, A. Amorosi, Sediment core analysis using artificial intelligence, *Sci. Rep.* 13 (1) (2023) 20409.
- [205] R. Achtabat, M. Dreyer, I. Eisenbraun, S. Bosse, T. Wiegand, W. Samek, S. Lapsuschkin, From “where” to “what”: Towards human-understandable explanations through concept relevance propagation, 2022, pp. 1–87, arXiv preprint [arXiv:2006.03208](https://arxiv.org/abs/2006.03208).
- [206] L. Yu, W. Xiang, J. Fang, Y.-P.P. Chen, L. Chi, eX-ViT: A novel explainable vision transformer for weakly supervised semantic segmentation, *Pattern Recognit.* 142 (2023) 109666.
- [207] Y. Huang, C. Qiu, K. Yuan, Surface defect saliency of magnetic tile, *Vis. Comput.* 36 (1) (2020) 85–96.
- [208] M.S. Bedmutha, S. Raman, Using class activations to investigate semantic segmentation, in: *Proceedings of the Computer Vision and Image Processing Conference*, Prayagraj, India, Revised Selected Papers, Part III 5, 2021, pp. 151–161.
- [209] X. Wu, Z. Li, C. Tao, X. Han, Y.-W. Chen, J. Yao, J. Zhang, Q. Sun, W. Li, Y. Liu, et al., DEA: Data-efficient augmentation for interpretable medical image segmentation, *Biomed. Signal Process. Control* 89 (2024) 105748.
- [210] V. Fischer, M.C. Kumar, J.H. Metzen, T. Brox, Adversarial examples for semantic image segmentation, 2017, pp. 1–4, arXiv preprint [arXiv:1703.01101](https://arxiv.org/abs/1703.01101).
- [211] M. Cisse, Y. Adi, N. Neverova, J. Keshet, Houdini: Fooling deep structured prediction models, 2017, pp. 1–12, arXiv preprint [arXiv:1707.05373](https://arxiv.org/abs/1707.05373).
- [212] C. Xie, J. Wang, Z. Zhang, Y. Zhou, L. Xie, A. Yuille, Adversarial examples for semantic segmentation and object detection, in: *Proceedings of the IEEE International Conference on Computer Vision*, 2017, pp. 1369–1378.
- [213] A.-K. Dombrowski, M. Alber, C. Anders, M. Ackermann, K.-R. Müller, P. Kessel, Explanations can be manipulated and geometry is to blame, in: *Proceedings of the Advances in Neural Information Processing Systems*, Vol. 32, 2019, pp. 1–12.
- [214] D. Hendrycks, K. Zhao, S. Basart, J. Steinhardt, D. Song, Natural adversarial examples, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 15262–15271.
- [215] Y.-N. Chuang, G. Wang, F. Yang, Z. Liu, X. Cai, M. Du, X. Hu, Efficient XAI techniques: A taxonomic survey, 2023, arXiv preprint [arXiv:2302.03225](https://arxiv.org/abs/2302.03225).
- [216] N. Cammarata, G. Goh, S. Carter, C. Voss, L. Schubert, C. Olah, Curve circuits, *Distill* (2021) <https://distill.pub/2020/circuits/curve-circuits>.
- [217] J. Marques-Silva, Logic-based explainability in machine learning, in: *Reasoning Web. Causality, Explanations and Declarative Knowledge: 18th International Summer School 2022*, Berlin, Germany, September 27–30, 2022, Tutorial Lectures, 2023, pp. 24–104.

- [218] K. Čyras, A. Rago, E. Albini, P. Baroni, F. Toni, Argumentative XAI: A survey, 2021, pp. 1–8, arXiv preprint [arXiv:2105.11266](https://arxiv.org/abs/2105.11266).
- [219] R. Karim, R.P. Wildes, Understanding video transformers for segmentation: A survey of application and interpretability, 2023, pp. 1–113, arXiv preprint [arXiv:2310.12296](https://arxiv.org/abs/2310.12296).
- [220] Z. Zhang, Z. Wang, I. Joe, CAM-NAS: An efficient and interpretable neural architecture search model based on class activation mapping, *Appl. Sci.* 13 (17) (2023) 9686.
- [221] Z. Carmichael, T. Moon, S.A. Jacobs, Learning interpretable models through multi-objective neural architecture search, 2021, pp. 1–25, arXiv preprint [arXiv:2112.08645](https://arxiv.org/abs/2112.08645).