

Contents lists available at ScienceDirect

Information Fusion



journal homepage: www.elsevier.com/locate/inffus

XEdgeAI: A human-centered industrial inspection framework with data-centric Explainable Edge AI approach

Hung Truong Thanh Nguyen^{a,*}, Loc Phuc Truong Nguyen^b, Hung Cao^a

^a Analytics Everywhere Lab, University of New Brunswick, Fredericton, E3B 5A3, New Brunswick, Canada
^b Friedrich-Alexander-Universität Erlangen-Nürnberg, Erlangen, 91054, Bavaria, Germany

ARTICLE INFO

ABSTRACT

Keywords: Explainable Edge AI Industrial Visual Quality Inspection Large Vision Language Model Recent advancements in deep learning have significantly improved visual quality inspection and predictive maintenance within industrial settings. However, deploying these technologies on low-resource edge devices poses substantial challenges due to their high computational demands and the inherent complexity of Explainable AI (XAI) methods. This paper addresses these challenges by introducing a novel XAI-integrated Visual Quality Inspection framework that optimizes the deployment of semantic segmentation models on lowresource edge devices. Our framework incorporates XAI and the Large Vision Language Model to deliver human-centered interpretability through visual and textual explanations to end-users. This is crucial for enduser trust and model interpretability. We outline a comprehensive methodology consisting of six fundamental modules: base model fine-tuning, XAI-based explanation generation, evaluation of XAI approaches, XAIguided data augmentation, development of an edge-compatible model, and the generation of understandable visual and textual explanations. Through XAI-guided data augmentation, the enhanced model incorporating domain expert knowledge with visual and textual explanations is successfully deployed on mobile devices to support end-users in real-world scenarios. Experimental results showcase the effectiveness of the proposed framework, with the mobile model achieving competitive accuracy while significantly reducing model size. This approach paves the way for the broader adoption of reliable and interpretable AI tools in critical industrial applications, where decisions must be both rapid and justifiable. Our code for this work can be found at https://github.com/Analytics-Everywhere-Lab/vgixai.

1. Introduction

Industrial Visual Quality Inspection systems are automated mechanisms, typically engineered to examine and continuously monitor the status of industrial hardware assets. In the rapidly evolving landscape of Artificial Intelligence (AI), especially Computer Vision (CV), the integration of vision techniques has revolutionized how manufacturers automatically maintain and ensure product quality and compliance, thereby mitigating human error and augmenting efficiency.

Deep Learning (DL)-based models, such as Deep Neural Networks (DNNs), have markedly improved the precision of numerous visual quality inspection systems has been significantly enhanced in terms of accuracy, efficiency, and running time. However, these advancements have significant trade-offs between accuracy, computational complexity, and interpretability. While achieving higher accuracy, the increased computational complexity of these models often results in a lack of interpretability. These opaque models, often perceived as "black boxes", pose challenges for domain experts and end-users in comprehending their internal decision-making processes. This opacity becomes a significant concern in sensitive domains, where decisions have profound implications [1–4].

The advent of XAI has steered a new era of model transparency and interpretability, fundamentally transforming how we understand and interact with AI systems. XAI methods enable domain experts to validate the model's reasoning process and identify potential biases or errors in the data [5–7] or model [8–12]. Despite the growing trend on XAI, there remains a significant gap in the practical application of these XAI techniques to enhance the performance and interpretability of visual quality inspection models, particularly in resourceconstrained environments like edge devices. This gap highlights a critical need for innovative solutions that adapt XAI methodologies for use in environments where computational resources are at a premium.

Our research introduces a cutting-edge framework for industrial visual quality inspection systems that integrates XAI and LVLM to (1) *improve the accuracy of semantic segmentation models for industrial*

* Corresponding author. *E-mail addresses:* hung.ntt@unb.ca (H.T.T. Nguyen), loc.pt.nguyen@fau.de (L.P.T. Nguyen), hcao3@unb.ca (H. Cao).

https://doi.org/10.1016/j.inffus.2024.102782

Received 7 June 2024; Received in revised form 25 October 2024; Accepted 2 November 2024 Available online 15 November 2024

^{1566-2535/© 2024} The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY-NC license (http://creativecommons.org/licenses/by-nc/4.0/).



Fig. 1. The XAI landscape, categorizing XAI techniques based on explanation scope, data type, modeling phase, architecture, and application domains.

assets and (2) provide visual and textual human-centered explanations to end-users. To this end, we address the following research questions:

- How can we effectively integrate XAI methods into Industrial Visual Quality Inspection systems to provide meaningful explanations for model predictions?
- 2. How can domain expert insights, informed by XAI explanations, be utilized for data augmentation strategies to improve the performance of semantic segmentation models?
- 3. What are the optimal techniques for adapting and optimizing the enhanced model for deployment on edge devices, ensuring that it remains both effective and efficient in resource-constrained environments?
- 4. How can we design the delivery mechanisms for human-centered explanations on edge devices, making them accessible and comprehensible to end-users without prior expertise in AI or XAI?

To address these research questions, we propose a framework that is organized into six modules: base model finetuning, XAI-based explanation generation, XAI evaluation, XAI-guided data augmentation, edge model development, and the generation of visual and textual explanations for the edge model. Our contributions are as follows:

- 1. We develop a comprehensive XAI-integrated Visual Quality Inspection framework that incorporates multiple XAI techniques directly into the model development process, enhancing both interpretability and performance.
- 2. We validate the effectiveness of XAI-guided data augmentation, demonstrating substantial improvements in the performance of semantic segmentation models for industrial asset inspection.
- 3. We introduce a tailored mobile optimization algorithm that employs pruning and quantization techniques, enabling efficient deployment of our advanced models on resource-constrained edge devices.
- 4. We extend the functionality of our framework by integrating LVLM to provide clear, contextual textual explanations to endusers on edge devices, making the system's decisions transparent and understandable.

The remainder of the paper is structured as follows: Section 2 presents the XAI landscape and recent development of XAI-based model improvement methods. Section 3 provides an overview of related work in semantic segmentation, XAI techniques, the industrial visual quality inspection, the application of XAI in industrial settings and the LVLM utilization in the human-centered XAI context. Section 4 presents our methodology, detailing each module of the XAI-integrated Visual Quality Inspection framework. Section 5 describes the implementation details, including model architecture, the procedure of XAI and LVLM integration, and the mobile optimization algorithm. Section 6 presents our first experiment with experimental results on an industrial assets dataset, evaluating the performance of the models in different stages, the effectiveness of XAI-guided data augmentation, and the demonstration of visual and textual explanations. In Section 7, we further extend our framework with the data-centric approach, where we evaluate the framework ability on a new industrial dataset while inheriting findings and models from the previous experiment. Finally, in Sections 8 and 9, we discuss and conclude the key findings and limitations from the experimental results and future mitigation research directions.

2. Background

This section offers a comprehensive overview of the XAI landscape and the key factors driving emerging research trends in XAI. Additionally, we summarize common XAI-based approaches for enhancing AI model performance (see Fig. 1).

2.1. Explainable Artificial Intelligence (XAI) landscape

In this section, we delve into the comprehensive landscape of XAI, focusing on the critical elements that shape its development and application. Specifically, we examine the Modeling Phase, Scope of Interpretability, Explanation Stage, Explanation Capability, and Architecture Approachability, as well as the significance of data-centric approaches in XAI. Additionally, we discuss various Application Domains to illustrate the practical impact of XAI across different fields. This detailed overview aims to provide a clear understanding of how XAI techniques are categorized and applied, shedding light on the diverse range of interpretability methods.

2.1.1. Modeling phase

The explainability of an AI model can be divided into three main phases: Pre-Modeling, Modeling, and Post-Modeling [13].

- **Pre-Modeling:** This phase focuses on analyzing, characterizing, and exploring the input data to understand the underlying properties of its features or to interpret the overall dataset summary. Typical techniques for data explanation include dataset descriptions and standardization methods (e.g., metadata descriptions, provenance, variable relationships, statistics, ground truth correlation, data cleaning, normalization, standardization), exploratory data analysis, explainable feature engineering, dataset summarization and enhancement (e.g., identifying meaningful abnormalities, summarizing interpretable outliers, selecting prototypes, and enhancing data interpretability). These techniques are generally applied before model creation.
- **Modeling:** This phase is dedicated to designing, developing, and implementing explainable model architectures and algorithms. The aim is to create new, understandable architectures and algorithms or to transform existing "black box" models into transparent ones. Approaches in the Modeling Phase include mixed explanation and prediction (e.g., multimodal explanations, rationalizing neural predictions), architectural modifications for explainability, hybrid explainable models, and regularization for explainability. These methods are typically employed during the model creation process.
- **Post-Modeling:** This phase aims to extract explanations and enhance the interpretability of output results. Techniques used for explaining outputs include macro-explanations, output visualizations, and interpreting targets at different levels of complexity (e.g., comparing functional (external, interpretative) explanations with mechanistic (internal, algorithmic) explanations). These methods are typically applied after the model has been created.

2.1.2. Scope of interpretability

Model interpretability often involves understanding the rationale behind a model's decision-making process. There are two levels of interpretability: global explanations and local explanations [14].

- Global explanations: Global explanations provide an overall understanding of the key features in an AI model. This is achieved by analyzing how these features impact model accuracy or their influence on the model's predictions. Such explanations are valuable for guiding policy decisions and testing hypotheses about feature importance.
- Local explanations: A local explanation provides insight into how an AI model arrived at a specific prediction or decision. They can answer questions such as "Why did the model produce this output for this input?" or "What if this feature had a different value?" Local explanations are essential for validating and refining the model's decisions on a case-by-case basis, particularly when unexpected results occur.

2.1.3. Explanation stage

Model explanations can be achieved in two primary ways: by adopting a transparent model (intrinsic) or by interpreting the model after it has been trained (post-hoc) [14].

- **Intrinsic:** Intrinsically explainable methods are designed for transparency, providing direct access to their decision-making logic. Examples include linear regression models, decision trees, and models that generate explanations during the learning process.
- **Post-hoc:** Post-hoc explanation methods provide insights into "black-box" models, such as DNNs, which are not directly interpretable.

2.1.4. Explanation capability

Model interpretability is achieved through various explanation techniques, each tailored to specific needs and contexts. Typically, these explanations assist users in five forms: text, visualizations, examples, concepts, and feature attributions [15].

- By Text: Text-based explanations aim to provide insights into model decisions by generating relevant words, phrases, or sentences that describe the model's reasoning. These explanations can also include symbols that represent the model's functioning, effectively mapping the model's rationale into a comprehensible form.
- By Visualization: Visual explanations utilize images or visualizations to interpret a model's decisions and behavior. Techniques such as saliency maps, feature highlighting, and dimensionalityreduced visualizations highlight the features and parts of a data instance that are most relevant to the model's decision.
- By Example: Example-based explanations clarify model decisions by providing representative instances or examples similar to the query instance. These methods help users understand model behavior by comparing decisions to known examples. The types of examples can include counterfactuals (instances that significantly change the model's prediction), adversarial examples (instances that can deceive models), prototypes (representative instances of the entire dataset), and influential instances (training instances that most impact the model's parameters and predictions).
- By Concept: Concept-based explanations enhance our understanding of a model's decision-making process by conveying the driving factors behind a prediction in terms of humanunderstandable concepts.
- By Feature Attributions: Feature attribution methods calculate relevance scores to identify and emphasize the specific features of the input data that have a significant impact on the model's output. This approach is commonly used in text and image data, where it assigns importance to specific words or regions within the images, thereby clarifying their influence on the model's predictions.

2.1.5. Architecture approachability

Post-hoc XAI methods can be categorized as model-specific or model-agnostic, depending on their applicability [14].

- **Model-specific:** Model-specific methods are limited to specific types of models. For example, interpreting regression weights in a linear model is a model-specific approach, as the interpretation is unique to that model type.
- **Model-agnostic:** Model-agnostic methods can be applied to any AI model and are used after the model has been trained (post hoc). These methods typically analyze the relationships between input and output pairs. By definition, they do not access the model's internals, such as weights or structural information.

2.1.6. Data-centric in XAI

Data is fundamental in developing AI systems. For many years, efforts have been concentrated on model-centric AI, primarily using fixed benchmark datasets. This focus has led developers to prioritize refining their models to enhance performance, placing a heavy reliance on data quality. However, this approach often overlooks the fact that data can be flawed and unreliable. As a result, there has been a recent shift towards enhancing the reliability, integrity, and quality of data, rather than focusing solely on model improvements.

In the data-centric AI approach, models are kept locked and unchanged during the experimentation and fine-tuning processes, while datasets are continuously modified, improved, and augmented. This method ensures that models perform accurately and remain unbiased across various data distributions and conditions. Enhancing data explainability helps the AI community build more robust systems and tackle complex real-world problems.

Zha et al. [16] summarize the tasks involved in developing data to simplify input and improve diversity, fairness, and understandability. The authors claim that this procedure potentially leads to better model performance, generalization, and robustness. In the field of XAI, most methods focus on three main data types: tabular data, images, and text. On the other hand, there are limited XAI methods for tackling realworld challenges with time-series or graph data [17]. Several examples of data-centric approaches in XAI can be found in Section 2.1.7.

2.1.7. Application domains

XAI has gradually been integrated into multiple sectors, such as automated transport, healthcare, finance, and education. These realworld implementations have demonstrated XAI's significant potential in enhancing model transparency, trust, reliability, and user acceptance. Table 1 provides a summary of the state-of-the-art XAI applications across these sectors.

- Manufacturing: XAI techniques have been extensively applied in manufacturing to support the adoption of DL models in safetycritical aspects. These include predictive maintenance [1], energy consumption optimization [9], visual inspection [5,27], and qualification [28]. By integrating XAI, the reliability and efficiency of manufacturing processes are significantly enhanced.
- **Transportation:** In the transportation sector, XAI is essential for ensuring safety, enhancing system transparency, and fostering user trust in autonomous vehicles. By providing clear explanations of autonomous decisions, XAI helps diagnose and rectify potential issues, contributing to safer and more reliable automated transport systems [4]. Studies have introduced various XAI models to improve human–vehicle interaction, explain decisions made by advanced driver-assistance systems, and enhance the reliability of autonomous driving systems by mimicking human decision-making processes [18,19].
- Finance: Financial institutions face increasing pressure to ensure their AI models comply with evolving regulations aimed at safeguarding consumer rights and maintaining market stability. AI models equipped with XAI capabilities can provide personalized advice, product recommendations, and risk assessments, thereby enhancing customer experience and building trust. Achieving a balance between interpretability and model performance is crucial for transparency in decision-making processes related to investments, credit scoring, and risk management [20, 23]. Moreover, XAI improves fraud detection mechanisms by offering clear explanations of fraud detection decisions, enabling financial institutions to respond swiftly and accurately [3].
- Healthcare: XAI significantly benefits healthcare, pharmacy, and bioinformatics by enhancing diagnostic accuracy, building trust, and ensuring the ethical use of AI technologies. Specifically, XAI applications not only improve diagnostic and therapeutic outcomes but also unravel complex biological data, revolutionizing patient care and medical research. Recently, XAI has further increased diagnostic accuracy by making AI-driven diagnoses more understandable and trustworthy for both clinicians and patients. An increasing number of advanced XAI methods are being developed and implemented to interpret AI decision-making processes, yielding promising results [2,22,23].
- Education: Integrating XAI into learning systems can significantly enhance the interpretability of student data analysis, offering educators and students valuable insights into personalized learning patterns, content suggestions, performance metrics, and customized feedback mechanisms [24–26]. By providing clear explanations of these personalized details, XAI empowers learners to take control of their learning process, fostering a more engaging and effective educational experience.

2.2. XAI-based model improvement methods

In the rapidly expanding industrial sector, developers require comprehensive information on how AI models process data to maintain accuracy and reliability [29-31]. This is not just about understanding model decisions; it involves leveraging that knowledge to improve different aspects of model performance. To meet this demand, XAI applications have extended beyond theoretical explanations and visual aids to provide practical insights that actively refine models. However, a significant gap remains despite the wealth of XAI research. Very few studies offer a systematic approach to closely examine and classify how XAI can enhance both model interpretability and functionality [32]. This can create significant challenges for developers who may not fully comprehend the working mechanisms of emerging model improvement techniques, potentially resulting in their misuse or misinterpretation [33]. To address this issue, Weber et al. [34] have proposed a framework that categorizes XAI-based techniques according to the component of the training loop they optimize, such as data, intermediate features, loss function, gradients, or model architecture. Below are detailed descriptions of the five main augmentation types and some corresponding notable methods.

2.2.1. Data augmentation

XAI-based data augmentation leverages explanations to reshape data structure. It takes the original dataset and its attributions as inputs, generating augmented data by creating new samples or adjusting existing data distribution. The main objective is to minimize bias and errors that result from using the original data structure. This type of augmentation is often implemented early in the model training process, particularly in the initial forward–backward phase, which influences all subsequent components [34]. Approaches that utilize data augmentation techniques for improving model performance include [5–7].

2.2.2. Feature augmentation

XAI-based feature augmentation utilizes explanations to adjust the model's feature representations. First, intermediate features from a specific layer and their attributions are taken as input. Then, XAI is employed to evaluate each feature's importance, allowing key features to be selectively scaled, masked, or transformed. Thereby, feature augmentation indirectly influences all higher-level feature representations and subsequent training components, including gradient updates and the final model. This holistic approach helps reduce potential biases and enhances the model's generalization, resulting in more accurate and reliable predictions [34]. Techniques that improve model performance through this type of augmentation are proposed in [23,35, 36].

2.2.3. Loss augmentation

XAI-based loss augmentation modifies the loss function using insights derived from local explanations to guide the training process. This type of augmentation can take different forms, such as adding regularization terms that adjust the model based on feature relevance or modifying the loss function to focus on challenging or underrepresented data points. These adjustments ensure that the model not only focuses on reducing the overall error but also addresses specific inaccuracies identified through explanations, leading to a more balanced and equitable performance across various conditions. Due to its working mechanism, loss augmentation indirectly impacts all components of the backward pass as well as the final model [34]. Notable XAI-based model improvement methods employing this technique include [37–39].

Table 1

An overview of XAI applications across various domains categorized by their modeling phase, the scope of interpretability, explanation stage, explanation capability, target audiences, and model improvement augmentation types.

Domain	Ref.	XAI contributions	Phase		Scope		Stage		Car	Capability			Audience			Augmentation						
											suc											
			Pre-Modeling	Modeling	Post-Modeling	Global	Local	Intrinsic	Post-hoc	Text	Visualization	Example	Concept	Feature attributio	AI Experts	Domain Experts	End-users	Data	Feature	Loss	Gradient	Model
uring	[9]	Use 2D SHAP-clustered explanations in automatic tuning model's hyperparameters to predict power consumption		•		•			•					•	•							•
Manufact	[1]	Present a self-explained AutoML with an interactive visualization module in the field of predictive maintenance	•			•	•	•					•	•	•							•
	[5]	Employ LRP to interpret and improve a DL model for classifying metal surface defects		•			•		•		•			•	•	•		•				
Transportation	[18]	Enhance the trust in driver emotion and distraction detection models			•		•		•		•				•							
	[4]	Produce explanations for actions in response to the state of action-inducing objects that could create a hazard		•			•	•		•	•				•	•	•			•		
	[19]	Use SHAP to interpret XGBoost, thereby analyzing factors that influence users' willingness to engage in ridesharing			•		•		•		•			•	•	•						
	[20]	Enhance the trust and performance of model in a complex financial decision-making context			•	•	•		•				•				•					
Finance	[3]	Use LIME-based explainers for fraud detection model on how each instance contributes to the final model output			•		•		•					•	•	•						
	[21]	Comparatively apply SHAP, LIME, GradCAM, Saliency Maps to explain credit scoring model predictions			•		•		•	•	•					•						
	[2]	Enhance the trust for the two-stage object detectors in the thyroid nodule diagnosis system		•		•		•		•			•	•	•							
Healthcare	[22]	Propose an intrinsic explainable algorithm to enable the data validation by the medical experts		•		•	•	•			•		•	•	•	•		•				
	[23]	Implement SHAP and LIME for enhancing interpretability in cardiovascular disease prediction			•		•		•		•			•		•			•			
Education	[24]	Provide insights into factors influencing school dropout predictions			•	•	•		•					•		•						
	[25]	Identify features important to the development of an interactive knowledge discovery tool			•	•	•					•		•		•						
	[26]	Study counterfactual explanations focusing on student performance prediction			•		•					•				•						

2.2.4. Gradient augmentation

XAI-based gradient augmentation involves two specific types: feature gradient augmentation and parameter gradient augmentation. Feature gradient augmentation adjusts gradients of intermediate features using scaling, masking, or transformation techniques based on their importance determined by XAI. On the other hand, parameter gradient augmentation targets gradients of specific model parameters and relies on XAI insights to fine-tune them precisely. Both methods occur during the training stage, with feature gradient augmentation impacting a more significant part of the network by modifying gradient flow across all layers below the target. In contrast, parameter gradient augmentation focuses its effects on specific layers, minimizing its broader network effects. By prioritizing critical gradients during backpropagation, XAI-based gradient augmentation effectively guides the learning process, leading to improved performance, faster convergence, and higher data efficiency [34]. Several methods utilize this augmentation type to enhance model performance, including [40–42].

2.2.5. Model augmentation

XAI-based model augmentation incorporates two main techniques: pruning and quantization. Pruning focuses on the model's architecture, using explanations to assess the importance of its parameters. It systematically removes less critical parameters to reduce the model's complexity and storage requirements while maintaining the overall performance. Quantization, on the other hand, aims to optimize the precision of the model's parameters. It leverages XAI to find and modify parameters that can tolerate lower precision without causing significant performance degradation. It is worth noticing that both methods are implemented post-training and generally have minimal influence on aspects like accuracy or robustness. However, they are essential for increasing the model's computational efficiency and speeding up its inference, making it more suitable for deployment in resourceconstrained environments [34]. A wide range of model improvement methods employing this augmentation approach can be found in [1,9– 12].

3. Related work

In this section, we examine the latest advancements in semantic segmentation and the application of XAI techniques to enhance model interpretability and transparency, particularly in the context of industrial visual quality inspection. We also explore the emerging paradigm of Explainable Edge AI and current developments in human-centered XAI, which uses LVLM to generate explanations that are accessible to users without a background in AI or XAI.

3.1. Semantic segmentation

Semantic segmentation is a process in CV where each pixel in an image is classified into one of several predefined categories, thereby segmenting the image into regions with distinct object identities. Unlike image classification, which assigns a single label to an entire image, semantic segmentation provides a detailed, pixel-level understanding of the scene. This section covers recent advancements in semantic segmentation models and local post-hoc XAI methods tailored for semantic segmentation.

3.1.1. Semantic segmentation models

DL backbones such as VGG [43], YOLO [44], ResNet [45], and MobileNet [46] have revolutionized visual quality inspection by offering robust feature extraction capabilities that are essential for segmentation tasks. These backbones form the foundation of semantic segmentation models, enabling precise identification and classification of objects within an image.

Semantic segmentation is a crucial tool for visual quality inspection systems, as it enables these systems to focus on critical parts of an image while ignoring irrelevant regions. Notable examples of semantic segmentation models include FCN [47], LRASPP [48], and DeepLabV3 [49]. These models represent substantial advancements in semantic segmentation due to their high performance and applicability on mobile devices.

DeepLabV3 [49], with its innovative Atrous Spatial Pyramid Pooling (ASPP) [50] module, significantly enhances semantic segmentation models by capturing objects at multiple scales and improving boundary delineation. The latest iteration, DeepLabV3Plus [51], incorporates an encoder–decoder structure to further refine object boundaries and details, demonstrating superior performance in various segmentation benchmarks. These continuous improvements in model architecture, efficiency, and accuracy have greatly advanced visual quality inspection systems. By enhancing the segmentation of small objects and intricate details, these developments enable more precise and reliable inspection processes across multiple industries.

3.1.2. Local post-hoc XAI methods for semantic segmentation

This section introduces local post-hoc XAI methods specifically designed for semantic segmentation tasks. It is worth noticing that methods initially tailored for the classification task can also be adapted to work with the outputs of semantic segmentation models [52]. These methods can be categorized based on their explanation generation mechanisms, including Backpropagation-based, Class Activation Mapping (CAM)-based, Perturbation-based [53], and Example-based methods (Fig. 2):



Fig. 2. Local post-hoc XAI methods for the semantic segmentation.

- **Backpropagation-based**: Backpropagation-based XAI methods explain neural network predictions by using the backpropagation algorithm to compute the gradients of the output with respect to the input features. These gradients reveal how changes in each input feature impact the output, thereby highlighting the most influential features [54,55].
- **Class Activation Mapping (CAM)-based:** CAM-based XAI methods explain neural network predictions by producing heatmaps highlighting important regions in the input data. They work by extracting feature maps from the final convolutional layer, weighting them according to their relevance to the predicted class, and summing the weighted feature maps to generate the heatmaps. This process identifies which parts of the input are most influential in the model's decision [56–65].
- **Perturbation-based**: Perturbation-based XAI methods explain neural network predictions by systematically modifying parts of the input data and observing changes in the model's output. This involves altering specific input features, such as occluding image regions, adding noise, or changing values in structured data. By comparing the original and perturbed outputs, these methods identify which changes significantly impact predictions, indicating the importance of the corresponding input parts. The results are used to create heatmaps highlighting the most influential regions or features [66–71].
- **Example-based**: Example-based XAI methods explain neural network predictions by identifying and presenting similar instances from the training data. They extract features from the input, measure similarity with training examples, and find the nearest neighbors. The most similar examples and their predictions are then shown to the user, helping them understand how the model makes decisions based on these similarities [72,73].

In this paper, we employ XAI methods that offer visual explanations, with a focus on CAM-based techniques such as GradCAM [58], GradCAM++ [57], XGradCAM [62], HiResCAM [74], ScoreCAM [59], AblationCAM [60], GradCAMElementWise [75], EigenCAM [61], and EigenGradCAM [61]. These techniques are chosen for their applicability and effectiveness in semantic segmentation tasks. Additionally, we incorporate a perturbation-based XAI method, namely RISE [68], due to its model-agnostic working mechanism.

3.2. XAI in industrial visual quality inspection

Visual quality inspection is crucial in manufacturing and related industries for ensuring product quality and reducing costs. Traditional methods, however, are often time-consuming and expensive [76]. Recent advancements in deep learning have revolutionized visual quality inspection across sectors such as automotive, electronics, and construction [77–79]. These AI-driven systems have demonstrated exceptional performance in automating complex inspection tasks, including steel bar quality assessment in construction [27], pear quality inspection in the food industry [80], and guided visual inspection for asset maintenance [28].

As AI becomes increasingly prevalent in industrial visual quality inspection, the importance of XAI has become more prominent. XAI aims to make AI systems more transparent and interpretable, fostering better collaboration between human experts and AI [15]. In the industrial quality control and inspection sector, several studies have proposed leveraging XAI to enhance model quality and transparency in visual inspection tasks. For example, Rovzanec et al. [81] introduced a framework where explanations provide feedback to inspectors, helping improve the underlying classification model for visual defect inspection. Lupi et al. [82] developed a framework for reconfigurable vision inspection systems that use XAI to adapt to varying product types and manufacturing conditions, making XAI more accessible to non-specialist users. Gunraj et al. [83] presented SolderNet, a deep learning-driven system for inspecting solder joints in electronics manufacturing, which incorporates XAI to explain its predictions and enhance trust.

Overall, integrating XAI into inspection systems greatly enhances transparency, adaptability, and user-friendliness, resulting in improved performance and collaboration between human experts and AI. As the field evolves, adopting XAI in industrial visual quality inspection is anticipated to drive more efficient and reliable quality control processes across various industries [84].

3.3. Edge explainable AI (XEdgeAI)

The fusion of XAI with edge computing, known as Edge Explainable AI (XEdgeAI), marks a significant paradigm shift with the potential to transform a wide range of industries. This innovation aims to make automated systems more transparent, understandable, and trustworthy.

Recent studies have illustrated the potential of XEdgeAI in various fields. For instance, Kok et al. [85] developed an XAI-powered edge computing solution for optimizing energy management in smart buildings. This system provides explanations of energy usage patterns and decision-making processes, enabling building managers to implement informed energy-saving changes. In another study, Garg et al. [86] addressed the challenges of building trust in AI systems within a 6G edge cloud environment, highlighting the critical role of transparency and interpretability in edge computing. In the healthcare domain, Dutta et al. [87] designed a human-centered XAI application for edge computing to ensure that healthcare professionals can understand and act on AI-driven decisions. This approach has great potential to enhance the quality and reliability of healthcare services.

While XEdgeAI has made significant progress, it still faces several challenges such as limited resources and security concerns. Therefore, it is essential to integrate XAI into both edge and cloud environments, as each plays a critical role in developing intelligent systems. As researchers continue exploring the potential of XEdgeAI, this field will likely shape the future of AI across various industries. Applications such as industrial predictive maintenance and quality inspection in edge computing environments are expected to benefit greatly from these advancements.

3.4. Human-centered XAI with LVLMs

Researchers have employed various innovative strategies to make AI systems more transparent and comprehensible to users without a background in AI or XAI. These strategies include aligning AI explanations with human psychology [88], simplifying algorithms [89], providing interactive explanations [90], and offering textual explanations [91–95].

Recent advancements in Large Language Model (LLM) have led to the development of LVLMs, which blends language understanding with vision encoding and reasoning. These models excel in tasks such as image captioning, document understanding, visual question answering, and multi-modal in-context learning [96–105]. This progress introduces new opportunities for integrating LVLMs to generate textual explanations for visual perception tasks, thereby enhancing explainability [106]. As a good example [107], X-IQE evaluates text-to-image generation methods by generating textual explanations using a hierarchical Chain of Thought (CoT) within MiniGPT-4 as the base LVLM. On the other hand, the LangXAI framework [106] demonstrates how integrating LVLMs can improve the understandability of visual AI systems. By generating textual explanations for saliency maps, LangXAI also emphasizes the importance of user-centric design.

4. Methodology

Building upon the advancements of segmentation for visual quality inspection, XAI-based model improvement, and human-centered explanations, we propose a novel XAI-integrated Visual Quality Inspection framework that enhances the interpretability and efficiency of visual quality inspection models on edge devices. Our framework integrates state-of-the-art semantic segmentation models, such as DeepLabv3Plus [51], with advanced XAI methods to generate highly plausible and faithful explanations for the model's decisions. Furthermore, we introduce an XAI-guided data augmentation module that leverages expert knowledge to improve the base model's performance, addressing the need for continuous model refinement in industrial settings. To ensure the framework's applicability in real-world scenarios, we focus on developing an efficient edge model that can be deployed on mobile devices, enabling on-site visual quality inspections. Additionally, we incorporate a human-centric explanation module that generates both saliency maps and textual explanations using LVLMs, such as GPT-4 Vision [103], to make the inspection results more accessible and understandable to end-users.

Our proposed XAI-integrated Visual Quality Inspection framework consists of six main modules as illustrated in Fig. 3: base model finetuning, base model explanation with XAI, XAI evaluation, XAI-guided data augmentation, edge model development, and saliency and textual explanation for the edge.

- 1. **Base Model Finetuning:** In the first module, we prepare the visual quality dataset and finetune a semantic segmentation model, DeepLabv3Plus, to serve as our base model. The dataset is preprocessed and split into training and validation sets. The model is trained using the Dice loss function, which is well-suited for imbalanced classes in image segmentation tasks.
- 2. Base Model Explanation with XAI: The second module focuses on explaining the base model using various XAI methods, including GradCAM, GradCAM++, XGradCAM, HiResCAM, Score-CAM, AblationCAM, GradCAMElementWise, EigenCAM, Eigen-GradCAM, and RISE. These methods generate saliency maps highlighting the regions in the input image that have the highest influence on the model's segmentation decision.
- 3. XAI Evaluation: In the third module, we evaluate the XAI methods using plausibility and faithfulness metrics. Plausibility is assessed using the Energy-Based Pointing Game (EBPG), Intersection over Union (IoU), and Bounding Box (Bbox), which measure how well the explanations align with human intuition. Faithfulness is evaluated using Deletion and Insertion metrics, which quantify the alignment between the explanations and the model's predictive behavior. Based on these evaluations, the most suitable XAI method is selected.
- 4. XAI-guided Data Augmentation: The fourth module involves XAI-guided data augmentation. The chosen XAI method is used to guide the annotation augmentation process, where the training set annotations are relabeled based on expert recommendations. The model is then retrained on the enhanced training dataset to demonstrate the potential of annotation augmentation in improving semantic segmentation models.



Fig. 3. The methodology of the XAI-integrated Visual Quality Inspection framework integrated with XAI methods with 6 building modules: (1) Base Model Finetuning with a provided visual quality dataset, (2) Base Model Explanation with XAI, (3) XAI Evaluation, (4) XAI-guided Data Augmentation to improve the base model performance, (5) Edge Model Development on mobile devices and (6) Saliency and Textual Explanation for the Edge. The end-users interact with the framework via a mobile application, while the domain experts can interact via a web application.



Fig. 4. The implementation details of our proposed XAI-integrated Visual Quality Inspection framework.

5. Edge Model Development: In the fifth module, we develop the edge model by quantizing and optimizing the improved base model for deployment on mobile devices. Dynamic quantization is applied to specific layers to reduce model size, and the model is converted to TorchScript format for efficient execution.

Additional optimizations are performed to enhance the model's performance on mobile devices.

6. Saliency and Textual Explanation for the Edge: Finally, in the sixth module, we generate visual explanations as saliency maps and textual explanations for the segmentation results obtained by the mobile model on edge devices. The chosen XAI method is used to generate the saliency maps. At the same time, a LVLM, such as GPT-4 Vision, is utilized to generate humanreadable textual explanations based on the segmented image and explanation map.

By following this methodology, our framework aims to provide a high-plausible, interpretable, and efficient visual quality inspection solution for industrial assets that can be effectively deployed on edge devices.

5. Implementation

This section presents the implementation details of our proposed XAI-integrated Visual Quality Inspection framework, comprising six main modules (see Fig. 4). We describe the data preparation, base model finetuning, XAI methods, evaluation metrics, XAI-guided data augmentation, edge model development, and the generation of visual and textual explanations for the mobile model on edge devices. The implementation leverages and evaluates state-of-the-art techniques, such as the DeepLabv3Plus model, various XAI methods, and the LVLM, to create an interpretable and efficient visual quality inspection system.

5.1. Module 1 – Base model finetuning

In this module, we present the process of preparing the visual quality dataset to build a base model by finetuning the semantic segmentation model.

5.1.1. Data preparation and preprocessing

The acquired dataset D contains the image set I and the corresponding annotation set A.

The varying sizes of images are handled by being dynamically adjusted, where image $I \in \mathcal{I}$ is processed through two primary transformations: conversion to a tensor and normalization. Initially, we convert an image from its native representation, where pixel values are in [0, 255], to a tensor format with values normalized to [0, 1], using $I'_{cij} = \frac{I_{cij}}{255}$ for each pixel I_{cij} in channel *c*. Subsequently, we apply channel-wise normalization to this tensor, adjusting each pixel value to zero mean and unit variance by the formula:

$$I_{cij}'' = \frac{\left(\frac{I_{cij}}{255}\right) - \mu_c}{\sigma_c} \tag{1}$$

where: $\mu = [0.485, 0.456, 0.406]$ are the mean values for the RGB channels, $\sigma = [0.229, 0.224, 0.225]$ are the standard deviation values for the RGB channels.

The corresponding Common Objects in Context (COCO) annotations are stored in JSON format, where the annotated masks are supported for the semantic segmentation task. These masks are grouped by object categories and generated by drawing polygons around the specified objects, and subsequently used as ground truth for model training. In detail, the mask $M_{ij} : \mathbb{R}^2 \to \{0, 1\}$ indicates the presence (1) or absence (0) of an object *j* in image *i*, based on the polygon coordinates provided in A_i .

We divide the original dataset D into an 80%–20% training D_{train} and validation D_{val} sets, with all images $I \in I$ resized to 700 × 700 pixels, with the resizing operation R(I), such that $R : \mathbb{R}^{H \times W \times C} \rightarrow \mathbb{R}^{700 \times 700 \times C}$. The training set D_{train} is used to finetune the model. In contrast, the validation set D_{val} is leveraged to extract the explanation for the model's prediction, evaluate the XAI methods' performance, and support the domain expert to identify any data error.

5.1.2. Model finetuning

In this section, we describe the finetuning process for the base visual quality inspection model Θ . The fine-tuning of the DeepLabv3Plus model is conducted using PyTorch.



Fig. 5. The architecture of the DeepLabv3Plus model, featuring an encoder with Atrous Spatial Pyramid Pooling and a decoder for upsampling and refining segmentation outputs.

Model preparation. The DeepLabv3Plus architecture is set as the base model Θ , where its architecture is illustrated in Fig. 5. This model combines the strengths of encoder–decoder architecture with atrous separable convolution to enhance segmentation precision, particularly along object boundaries.

The encoder part of DeepLabv3Plus adopts a deep Convolutional Neural Network (CNN) model as the backbone, such as ResNet [45], or MobileNet [46], which is modified for segmentation tasks by incorporating depthwise separable convolutions that reduce the number of parameters and computational costs while maintaining effective feature extraction. The ImageNet pre-trained weights [108] are used for encoder initialization. The atrous convolution allows the network to grasp contextual information at multiple scales through the ASPP module without increasing computational demand.

The decoder module focuses on refining the segmentation outputs, which is critical for achieving high-resolution and accurate boundary representation. By gradually restoring spatial information and enhancing feature resolution, the decoder aids in producing more precise segmentation maps. The weights and activations of the convolution layers are in the Single-precision floating-point format (float32).

Loss function and optimization. The Dice loss function is used for training the model, which is particularly useful for imbalanced classes in the image segmentation task, as it considers the overlap between the predicted and ground truth masks [109]. Let P and GT represent predicted and ground truth masks, respectively. The Dice loss $\mathcal{L}_{\text{Dice}}$ is defined as:

$$\mathcal{L}_{\text{Dice}}(\mathbf{P}, \mathbf{GT}) = 1 - \frac{2 \times |\mathbf{P} \cap \mathbf{GT}|}{|\mathbf{P}| + |\mathbf{GT}|}$$
(2)

The accuracy of segmentation models is assessed using the IoU metric, also known as the Jaccard coefficient [110], one of the most commonly used metrics in semantic segmentation. The IoU metric is a prevalent evaluation measure in semantic segmentation. It quantifies the overlap between the predicted segmentation and the ground truth for each class. Mathematically, for the predicted mask P_c and the ground truth mask GT_c of a category, the IoU is calculated as follows:

$$IoU_c = \frac{P_c + IGI_c}{P_c \cup GT_c}$$
(3)

The model is optimized using the Adam optimizer [111] with an initial learning rate $\alpha = 0.0001$. α is halved to 1e–5 after 25 epochs to facilitate convergence as the model approaches optimal performance. The batch size is set to 8, balancing computational efficiency and memory constraints during training.

5.2. Module 2 - Base model explanation with XAI

We implement several XAI methods in this module to explain the semantic segmentation model. The explanation maps of all methods are extracted from the predictions of the segmentation model on the validation set, which will be used for the evaluation step. We utilize several notable CAM-based XAI methods, such as GradCAM [58], Grad-CAM++ [57], XGradCAM [62], HiResCAM [74], ScoreCAM [59], Ab-lationCAM [60], GradCAMElementWise [75], EigenCAM [61], Eigen-GradCAM [61] due to their applicabilities and plausibility in the semantic segmentation task. Besides, we also leverage a perturbation-based XAI method, namely RISE [68], due to its model-agnostic mechanism.

The resulting saliency map L^c of a specific category c highlights the regions in the input image with the highest influence on the model's segmentation decision. The saliency map is typically represented as a heatmap, where warmer colors indicate regions of higher importance, and cooler colors represent regions of lower importance.

5.3. Module 3 – XAI evaluation

This component evaluates the XAI methods with plausibility and faithfulness metrics on their explanations of the models with the validation set D_{val} . Plausibility measures how well the explanations align with human intuition and understanding, while faithfulness measures how accurately the explanations reflect the underlying model's decision-making process. By evaluating both plausibility and faithfulness, we can ensure that the chosen XAI method provides explanations that are both understandable to humans and accurately represent the model's behavior. Eventually, the method achieving the highest scores in most metrics will be chosen as the core XAI method of the model enhancement step. In the following, we introduce two relevant metrics, including the plausibility and faithfulness of XAI explanations.

5.3.1. Plausibility evaluation metrics

Plausibility, the alignment of explanations with human intuition, is assessed using measures like Energy-Based Pointing Game (EBPG) [59], and Intersection over Union (IoU) [56], and Bounding Box (Bbox) [112]. Based on human annotations, these measures validate the model by assessing the statistical superiority of explanations.

• Energy-Based Pointing Game (EBPG): evaluates the precision and denoising ability of XAI methods to identify the most influential region in an image for a given prediction [59]. It calculates how much the energy of the saliency map by pixel $L^{c}(i, j)$ falls inside the ground truth. A good explanation is considered to have a higher EBPG. EBPG formula is defined as follows:

$$EBPG = \frac{\sum L_{(i,j)\in GT}^c}{L_{(i,j)\notin GT}^c + L_{(i,j)\notin GT}^c}$$
(4)

• **IoU** [56] assesses the localization capability and the significance of the attributions captured in an explanation map. It measures the overlap between the saliency map and the ground truth annotation. IoU is defined as:

$$IoU = \frac{Area(L^c \cap GT)}{Area(L^c \cup GT)}$$
(5)

• **Bounding Box (Bbox)** [112] is a variant of the IoU metric that adapts to the size of the object of interest. It measures the overlap between the bounding box of the saliency map and the ground truth bounding box. Bbox is defined as:

$$Bbox = \frac{Area(BBox(L^{c}) \cap BBox(GT))}{Area(BBox(L^{c}) \cup BBox(GT))}$$
(6)

5.3.2. Faithfulness evaluation metrics

Faithfulness, the alignment of explanations with the model's predictive behavior, is evaluated using the Deletion and Insertion metrics [68,113,114]. These measures quantify the degree to which the explanations align with the predictive behavior of the model.

Given a black box model Θ , input image *I*, saliency map L^c , and number of pixels *N* removed per step.

• **Deletion** measures the accuracy of saliency areas by removing pixels from the input image in order of saliency, from large to small. More accurate saliency areas will have a steeper deletion curve, and a smaller deletion metric value indicates higher accuracy. Deletion d_{Del} is defined as:

$$d_{\text{Del}} = \text{AreaUnderCurve}(h_i \text{ vs. } i/n, \forall i = 0, \dots, n)$$
(7)

where $h_i \leftarrow f(I)$ while *I* has non-zero pixels, and according to L^c , setting the next *N* pixels in *I* to 0 each iteration until *n* iterations.

• **Insertion** assesses the comprehensiveness of the saliency area by removing all pixels from the input image and then recovering them in order of saliency, from large to small. A more comprehensive saliency area will require fewer pixels to recover the object and have a faster-rising insertion curve. A higher insertion metric value indicates a more comprehensive saliency area. The Insertion *d*_{Ins} is defined as:

$$d_{\text{Ins}} = \text{AreaUnderCurve}(h_i \text{ vs. } i/n, \forall i = 0, \dots, n)$$
(8)

where $I' \leftarrow \text{Blur}(I)$ and $h_i \leftarrow f(I)$ while $I \neq I'$, and according to L^c , setting the next *N* pixels in *I'* to the corresponding pixels in *I* each iteration until *n* iterations.

Finally, we choose the most suitable XAI method \mathcal{X} based on the evaluation results with the validation set \mathcal{D}_{val} and their capabilities for running on the mobile model θ .

5.4. Module 4 - XAI-guided data augmentation

Data augmentation strategies, such as altering data distribution or adjusting data and labels, have been employed to enhance model performance [115]. In this module, we leverage the advisable XAI method \mathcal{X} , demonstrating the highest faithfulness and plausibility from the XAI evaluation step, to guide the annotation augmentation process.

To facilitate the augmentation process, we develop a web-based user interface that allows domain experts to load any sample from the dataset and monitor the model's predictions, explanations in the form of saliency maps, and textual explanations. Fig. 6 illustrates the user interface, which consists of several components. The input image is displayed on the left, while the segmentation output is shown in the center. The saliency map, generated by the advisable XAI method \mathcal{X} , is presented on the right, highlighting the regions that contribute most to the model's prediction. At the bottom, a textual explanation is provided, offering a human-readable interpretation of the model's decision-making process. The web-based user interface plays a crucial role in this module, enabling domain experts to closely examine the model's behavior and provide informed recommendations for annotation augmentation.

Using this web UI, domain experts can thoroughly examine the model's predictions and explanations, identifying problems on the A_{val} and defining solutions for A_{train} need to be refined. Based on their expertise and the insights gained from the explanations, the experts provide recommendations for relabeling the annotations. These recommendations are then used to augment the training set, resulting in an enhanced dataset with improved annotations \tilde{A}_{train} .

Subsequently, the base model Θ is retrained on the augmented training dataset $\tilde{\mathcal{A}}_{\text{train}}$ to achieve the enhanced model $\tilde{\Theta}$. To assess the impact of the XAI-guided data augmentation, we evaluate the enhanced model $\tilde{\Theta}$ on the original validation set Dval by comparing the performance of the base model Θ before and after applying the data augmentation. Hence, we can demonstrate the potential of annotation augmentation, supported by XAI explanations, in enhancing semantic segmentation models.



Fig. 6. Web-based user interface for domain experts to monitor predictions, explanations, and textual explanations. The input image, segmentation output, saliency map, and textual explanation are displayed, allowing experts to assess the model's performance and provide feedback for annotation augmentation.

5.5. Module 5 – Edge model development

In this subsection, we will explore the process of transferring the original model to edge devices. Given the limited computational resources of mobile appliances, adaptations are needed to ensure smooth compatibility. After optimization, the model undergoes various developments to be deployed on platforms like Android and iOS. Detailed descriptions of these processes are presented in the subsequent parts.

5.5.1. Model quantization, pruning and optimization for mobile devices

After acquiring the enhanced model $\tilde{\Theta}$, we apply quantization, pruning, and optimization techniques to convert it into a mobile model θ that can efficiently run on smartphone devices. Algorithm 1 summarizes the entire model quantization, pruning, and optimization process for mobile deployment.

The first step is to disable batch normalization layers in the base model. We iterate over all the modules in the model and set the batch normalization layers to evaluation mode. This step is necessary to ensure that the model's statistics remain fixed during quantization. Next, we apply dynamic quantization to the base model Θ . Dynamic quantization is a technique that reduces the numerical precision of the model's weights and activations, thereby decreasing the model size and improving inference speed without significantly compromising accuracy. We target specific layers that are known to consume a substantial amount of computational resources, including the 2D convolutional layer (Conv2d), linear layer (Linear), rectified linear unit (ReLU), 2D batch normalization (BatchNorm2d), and 2D adaptive average pooling (AdaptiveAvgPool2d). We considerably reduce the model size by quantizing these layers' parameters to 8-bit integers using the dynamic quantization function from PyTorch, making it more suitable for deployment on devices with limited storage and processing capabilities. We perform a forward pass on the model using a sample input tensor to simulate the inference process and ensure the quantized model functions correctly.

After quantization, we apply pruning to the model. Pruning is a technique that removes less important weights from the model, reducing its size and computational requirements. We iterate over the named modules of the base model and apply structured pruning to the convolutional layers (Conv2d). We set the pruning amount to 0.1, indicating that 10% of channels in each convolutional layer will be pruned. To remove the pruning re-parameterizations and obtain the final pruned model, we iterate over the named modules again and remove the pruning masks.

Next, we trace the pruned model using the just-in-time (JIT) compiled tracing function, which converts the PyTorch model into a Torch-Script representation. TorchScript is a static graph representation that allows for optimizations and efficient execution on various platforms, including mobile devices. We pass a sample input tensor to the tracing function to capture the model's computational graph.

Finally, we apply mobile-specific optimizations to the traced model using the mobile optimization function. This function performs a series of optimizations tailored for mobile environments, such as operator fusion, constant folding, and dead code elimination. These optimizations help to reduce the model size further and improve its inference speed on mobile devices.

The resulting mobile model θ is now quantized, pruned, and optimized for deployment on smartphone devices. The quantization process reduces the model's memory footprint, the pruning process removes redundant weights, and the mobile-specific optimizations enhance the model's efficiency during inference.

5.5.2. Mobile model deployment

After acquiring the optimized mobile model θ , we proceed to the deployment process of the smartphone application, which incorporates the mobile model θ for both Android and iOS platforms. The deployment workflow encompasses packaging the model and its dependencies and integrating them with the mobile app. For Android devices, we leverage Maven, which is the process of compiling the app, bundling the mobile model, and generating an Android Package (APK) or Android App Bundle (AAB) for distribution, which is automated. While for iOS devices, we utilize CocoaPods, a dependency manager for Swift and Objective-C projects. Subsequently, the app is built and packaged using Xcode's build system.

Alg	gorithm 1: Model Quantization, Pruning, and Optimization								
for	Mobile								
	Input: base_model Θ : a base model								
	Input: input_tensor: an input tensor								
	Output: mobile_model θ : an optimized, quantized, and pruned								
	model	,							
	/* Disable batch normalization */								
1	for module in base_model.modules() do								
2	if isinstance(module, nn.BatchNorm2d) then								
3	module.eval()								
4	end								
5	end								
	/* Apply dynamic quantization */								
6	<pre>base_model = torch.quantization.quantize_dynamic(base_model,</pre>								
	{nn.Conv2d, nn.Linear, nn.ReLU, nn.BatchNorm2d,								
	nn.AdaptiveAvgPool2d}, dtype=torch.qint8)								
	/* Forward pass to simulate inference */								
7	base_model(input_tensor)								
	/* Apply pruning */								
8	for name, module in base_model.named_modules() do								
9	if isinstance(module, nn.Conv2d) then								
10	prune.ln_structured(module, name='weight',								
	amount=0.1, n=1, dim=0)								
11	end								
12	end								
	/* Remove pruning reparameterizations */	,							
13	for name, module in base_model.named_modules() do								
14	if isinstance(module, nn.Conv2d) then								
15	prune.remove(module, 'weight')								
16	end								
17	end								
	/* Trace the model */	,							
18	<pre>traced_model = torch.jit.trace(base_model, input_tensor)</pre>								
	<pre>/* Optimize the traced model for mobile */</pre>	,							
19	<pre>9 mobile_model = optimize_for_mobile(traced_model)</pre>								
20	o return mobile_model								

Once the mobile app is installed on the end users' devices, field engineers can utilize it to capture real-world images and request semantic segmentation using the integrated mobile model θ . The app provides an interface that guides field engineers to capture images, which leverages the device's camera capabilities. During the inference process, the mobile model generates a segmentation mask overlaying on the uploaded picture \tilde{I} as the prediction on the edge $y_{\tilde{I}}^{\theta}$ that identifies the visual quality objects being inspected in the image.

5.6. Module 6 – Saliency and textual explanation for the edge

In this section, we present the process of generating visual explanations as saliency maps and textual explanations for the segmentation results obtained by the mobile model θ on edge devices.

The visual explanations provide insights into the model's decisionmaking process, highlighting the important regions in the input image that contribute to the segmentation output. To generate the saliency map for the edge, the app first performs a forward pass of the input image through the mobile model θ to obtain the segmentation output. The chosen XAI method \mathcal{X} receives the uploaded image \tilde{I} , segmentation output $y_{\tilde{I}}^{\theta}$, the detected category *c* and interaction with the model θ to generate the explanation map.

The textual explanations, generated using LVLMs, offer a humanreadable interpretation of the segmentation results, enhancing the interpretability and trustworthiness of the mobile app for field engineers. In this research, we employ a recent member of the LVLMs family, GPT-4 Vision [103], as the core vision language model due to its robust Table 2

Summary of object characteristics by class in the TTPLA dataset, showing the number of objects, area percentages, and dimensions (width and height in pixels) with minimum, maximum, and average values.

	0			
Class	*#Obj.	Min/Max/Avg area (%)	Min/Max/Avg width (px)	Min/Max/Avg height (px)
cable tower_lattice tower_wooden tower_tucohy	10 082 404 333 232	0.00/9.88/0.26 0.04/58.28/7.22 0.01/43.61/2.46 0.02/23.68/3.25	4/3840/1181 34/3840/875 25/3684/484 29/3385/738	4/2160/766 59/2160/1226 112/2160/1061 137/2160/1322

performance across diverse tasks [116] and a strong alignment with human evaluators [117]. This LVLM processes a designed prompt, the segmented image, and the explanation map utilizing its pre-trained knowledge to generate human-readable explanations. The model's responses are based on its understanding of the visual content and its ability to associate relevant textual descriptions. The generated explanations provide a concise and intuitive summary of the segmentation results, highlighting the main objects the model concentrates on, their attributes, and their relationships within the image.

Finally, the textual explanations are displayed to the field engineer alongside the segmented image and the saliency map. This combination of visual and textual explanations enhances the interpretability of the segmentation results, allowing field engineers to understand the model's behavior better and the rationale behind its predictions.

Fig. 7 illustrates the iOS user interface of the mobile application designed for iPhone 11 Pro. The main screen allows users to upload an image by selecting from the photo gallery or capturing a new image using the device's camera. Upon uploading an image, the user can initiate the visual quality inspection process by choosing the object's category being inspected and tapping the "Inspect" button. The mobile model θ then performs semantic segmentation on the uploaded image \tilde{I} , generating a segmentation mask $y_{\tilde{I}}^{\theta}$ that identifies the visual quality objects being inspected. The segmentation mask is overlaid on the original image, providing a clear visual representation of the detected objects.

6. Experiment 1: A comprehensive evaluation

This section details the experimental setup, results, and analysis of our XAI-integrated Visual Quality Inspection framework, which is applied to an industrial hardware assets dataset for inspecting transmission towers and power lines using aerial imagery. We begin by training the base DeepLabv3Plus model with different backbones and assessing its performance on the validation set. Next, we conduct a series of comprehensive analyses to identify the optimal explaining method for implementation. Using this method, we apply data augmentation techniques to enhance the model's performance and generate textual explanations that provide human-understandable insights into the model's decision-making process. During the performance comparison stage, we consider the base model alongside the enhanced and mobile models.

6.1. Dataset

This experiment uses the TTPLA dataset for detecting and segmenting power-grid hardware components from aerial imagery [118]. The dataset encompasses 1242 high-resolution aerial images featuring 8987 instances of transmission towers and power lines. These instances are classified into four distinct categories: *cable, tower_wooden, tower_lattice,* and *tower_tucohy*, each representing a specific type of power-grid infrastructure component. Table 2 summarizes the object characteristics by class, detailing counts, area percentages, and dimension ranges (width and height) for analysis.

The dataset is annotated in the COCO format, facilitating detection and segmentation tasks, including semantic and instance segmentation.



Fig. 7. The iOS user interface of the mobile application for end-users, designed for iPhone 11 Pro. Users can upload an image, initiate visual quality inspection for a particular category, view the segmentation results along with explanations, and have the option to upload another image.



Fig. 8. Samples from the TTPLA dataset represent the main objects of categories in the green masks (a) *cable*, (b) *tower_wooden*, (c) *tower_lattice*, (d) *tower_tucohy*.

The diverse representation of objects against varying backgrounds, under different lighting conditions, and at multiple scales poses unique challenges to object detection and segmentation efforts. Furthermore, the TTPLA dataset supports the detection and semantic segmentation tasks and extends its utility to instance segmentation. This capability is crucial for identifying and differentiating between individual transmission towers and power lines, allowing for a deep analysis of the power grid infrastructure.

Fig. 8 showcases a selection of aerial images from the TTPLA dataset, illustrating the diversity and complexity of the power-grid infrastructure components it encompasses, highlighting the challenges involved in the segmentation task due to the diverse backgrounds, lighting conditions, and object sizes.

6.2. Base model performance

In this section, we evaluate the performance of the base model Θ after the finetuning process. We leverage the DeepLabv3Plus model with three different backbones, namely MobileNetv2, ResNet50, and ResNet101. Each model is trained on the D_{train} for 1000 epochs, where the training loss and accuracy are presented in Fig. 9. The training loss, shown in Fig. 9(a), steadily decreases for all three models, indicating successful learning and convergence during the finetuning process. The DeepLabv3Plus with ResNet101 backbone achieves the lowest training loss, followed by ResNet50 and MobileNetv2.

The results reported in Table 3 demonstrate the effectiveness of the DeepLabv3Plus architecture for the visual quality inspection on the validation set D_{val} . The ResNet101 backbone, with its deeper network structure, captures more complex features and achieves the highest



Fig. 9. The training loss in Dice loss and accuracy in mIoU(%) of three base DeepLabv3Plus models over 1000 epochs.

segmentation accuracy. The ResNet50 backbone balances performance and computational efficiency, while the MobileNetv2 backbone offers a lightweight option suitable for resource-constrained environments.

Overall, the finetuning process successfully adapts the base model Θ to the specific visual quality inspection task, achieving high IoU scores across different object categories, which can be a strong foundation for further optimization and deployment on edge devices.

6.3. Explanation evaluation

To select the most advisable XAI method \mathcal{X} for our framework, we perform both qualitative and quantitative evaluations of the explanations generated by several implemented XAI methods.

6.3.1. Qualitative evaluation

The explanation maps of implemented XAI methods for the base model Θ on the TTPLA validation set D_{val} are demonstrated in Fig. 10. The figure presents a visual comparison of the explanation maps generated by various XAI methods, including RISE, EigenGradCAM, Eigen-CAM, GradCAM, AblationCAM, GradCAMElementWise, GradCAM++, HiResCAM, ScoreCAM, and XGradCAM. The first row of the figure shows the input image, ground truth annotation, and the segmentation output produced by the model. The subsequent rows display the explanation maps generated by each XAI method.

Table 3

Accuracy comparison of DeepLabv3Plus (DLv3P) variants with different backbones (MobileNetv2, ResNet50, ResNet101) in terms of average IoU (%) for each category and mIoU (%). The models are evaluated at different stages: base (B), enhanced (E), and mobile (M) on the TTPLA validation set.

Model	#Params	Size (MB)	cable	tower_wooden	tower_lattice	tower_tucohy	mIoU (%)
DLv3P-MobileNetv2-B	4.37M	16.71	53.94	80.11	88.19	86.49	77.18
DLv3P-MobileNetv2-E	4.37M	16.71	54.37	86.49	80.98	88.78	77.82
DLv3P-MobileNetv2-M	3.51M	13.39	48.74	84.54	78.55	86.12	75.48
DLv3P-ResNet50-B	26.67M	101.76	56.66	92.31	93.18	90.63	83.20
DLv3P-ResNet50-E	26.67M	101.76	57.42	92.97	91.34	93.87	83.90
DLv3P-ResNet50-M	21.36M	81.48	53.67	90.39	88.80	91.26	81.53
DLv3P-ResNet101-B	45.66M	174.21	57.22	95.23	96.02	91.42	84.97
DLv3P-ResNet101-E	45.66M	174.21	58.36	96.54	92.89	97.61	86.35
DLv3P-ResNet101-M	36.57M	139.52	54.43	94.07	90.22	95.09	83.95



Fig. 10. The qualitative evaluation of XAI methods in explaining the base DeepLabv3Plus-ResNet101 model on a validation sample. The category is the tower_lattice. The IoU value between the segmentation and the ground truth is 96.25%.

From a qualitative perspective, the explanation maps should provide valuable insights into the model's decision-making process. We observe variations in the highlighted regions by comparing the explanation maps across different XAI methods. Some methods, such as RISE and EigenGradCAM, produce more localized and fine-grained explanations, accurately capturing the relevant objects and their boundaries. Other methods, like GradCAM and GradCAM++, generate more coarsegrained explanations, highlighting larger regions of interest. The qualitative evaluation also reveals the strengths and limitations of each XAI method. For instance, HiResCAM and ScoreCAM produce explanations with higher spatial resolution, enabling more precise localization of important features. On the other hand, methods like GradCAMElementWise and XGradCAM generate explanations with varying sensitivity levels to different image regions.

By visually comparing the explanation maps with the ground truth annotations and segmentation outputs, we can assess the plausibility of the explanations where plausible explanations should align well with human intuition and highlight regions that are semantically relevant to the target objects.

6.3.2. Quantitative evaluation

In this section, we present a quantitative evaluation of different XAI methods introduced in Module 2 (Section 5.2) on the validation set D_{val} using various plausibility and faithfulness metrics. The plausibility metrics include Energy-Based Pointing Game (EBPG), Intersection over Union (IoU), and Bounding Box (Bbox). In contrast, the faithfulness metrics include Deletion (Del) and Insertion (Ins), which are introduced in Module 3 (Section 5.3). The evaluation results are summarized in Table 4.

Among the evaluated methods, RISE stood out as the most advisable XAI method for our framework. RISE demonstrated exceptional performance in terms of faithfulness, achieving the best scores in both Deletion (0.123) and Insertion (0.691) metrics. These faithfulness results indicate that the explanations generated by RISE closely align with the model's predictive behavior, accurately capturing the most influential regions in the input images. It also demonstrated strong plausibility, with the second-best scores in EPBG (62.42%) and IoU (56.13%), and the best score in BBox (63.52%). RISE's high scores in these plausibility metrics demonstrate that its explanations are highly

Table 4

The quantitative result of XAI methods on the validation set \mathcal{D}_{val} on different metrics such as EPBG(%), BBox(%), IoU(%), Drop, Increase, and running time in seconds. For each metric, the arrow \uparrow/\downarrow indicates higher/lower scores are better. The best is in **bold**, and the second best is in <u>underline</u>.

Method	EPBG↑	BBox↑	IoU↑	Del↓	Ins↑
GradCAM	50.49	48.39	47.94	0.521	0.527
GradCAM++	58.13	52.24	53.22	0.517	0.547
HiResCAM	60.81	41.69	52.19	0.501	0.559
ScoreCAM	54.01	43.95	51.94	0.434	0.609
XGradCAM	57.94	47.81	53.09	0.594	0.551
AblationCAM	61.03	51.39	54.73	0.498	0.589
GradCAMElementWise	40.95	35.91	39.25	0.859	0.415
EigenCAM	51.49	42.54	51.03	0.545	0.491
EigenGradCAM	64.11	62.24	60.93	0.520	0.534
RISE	62.42	63.52	56.13	0.123	0.691

interpretable and closely match human understanding of the important regions in the input images.

EigenGradCAM also shows strong performance in plausibility metrics, achieving the highest EBPG and IoU scores of 64.11% and 60.93%, respectively. It also obtains the second-highest Bbox score of 62.24%. However, its faithfulness scores are not as impressive as RISE, suggesting that the salient regions identified by EigenGradCAM may not fully align with the model's predictive behavior. Other XAI methods, such as GradCAM++, XGradCAM, and AblationCAM, demonstrate competent performance in plausibility metrics but fall short in faithfulness compared to RISE.

Another advantage of RISE, particularly suitable for our visual quality framework, is its model-agnostic nature. As a model-agnostic XAI method, by treating the model as a "black box", RISE can be applied to any model without requiring access to its internal architecture or gradients. This property is crucial in our framework, where we may need to quantize or optimize the base model to a mobile model without being able to access or modify its architecture. RISE's flexibility ensures that we can still generate meaningful explanations for the model's predictions, regardless of any modifications made during the optimization process.

Based on our quantitative evaluation and considering the modelagnostic property of RISE, we use RISE as the advisable XAI method \mathcal{X} for our visual quality framework. The selection of RISE as the advisable XAI method has significant implications for the subsequent modules in the framework, such as Module 4 (Section 5.4) and Module 6 (Section 5.6). By incorporating RISE into our framework, we can enhance the interpretability and performance of our models, enabling the domain expert to validate the model's predictions while allowing end-users to understand the model's rationale even with the mobile model θ .

6.4. Model improvement via XAI-guided data augmentation

This section presents the experimental results of enhancing the DeepLabv3Plus-ResNet101 model's performance using annotation augmentation guided by the advisable XAI method \mathcal{X} and the domain expert. The process begins with the XAI method \mathcal{X} generating explanations for each image in the validation set \mathcal{D}_{val} . The domain expert, knowledgeable in semantic segmentation models and XAI algorithms, analyzes the saliency maps to guide data augmentation.

As shown in Fig. 11, the model effectively segments the cable from a clean or mixed-objects background, such as Fig. 11a. However, the model's performance decreases when the background contains objects resembling the target object, as shown in Fig. 11b. The explanations reveal that the model's attention is directed at the object and the surrounding background. However, the model lacks contextual attention to surrounding objects and backgrounds in complex cases. This behavior is due to the ability of models to leverage local and global



Fig. 11. List of input images, ground truth, segmentation of the base DeepLabv3Plus-ResNet101-B model for the *cable* inspection for two segmentation cases: (a) Good segmentation and (b) Bad segmentation.



(a) Annotation Enlargement



(b) Adding Annotations for Perplexed Objects

Fig. 12. Annotation augmentation approaches: (a) Annotation enlargement where the annotation size for thin objects like cables is increased by 2 pixels on both sides. (b) Adding annotations for perplexed objects like the road surface marks to guide the model in differentiating between white cables and perplexed objects.

contextual information from the original annotations [69]. A domain expert suggests annotation augmentation for each sample to enhance the model's performance.

Two approaches are proposed, namely *Annotation Enlargement* and *Adding Annotations for Perplexed Objects*, as illustrated in Fig. 12, which are described as follows:

- Annotation Enlargement (Fig. 12(a)): Given that the model can leverage surrounding contextual information to improve its performance, we propose to enlarge the annotations of thin objects, especially thin cables, which the model often overlooks based on the saliency maps. We increase the object's size by 2 pixels on both sides.
- Adding Annotations for Perplexed Objects (Fig. 12(b)): As the model often confuses cables with perplexing objects like road surface markings, we propose adding *void* annotations to categorize these perplexing objects as unlabeled objects.

The enhanced DeepLabv3Plus-ResNet101 (DLv3P-ResNet101-E) model demonstrates improved segmentation of thin objects from the background and perplexing objects (Fig. 13). The IoU of the enhanced model $\tilde{\Theta}$ is also higher than that of the base version Θ , particularly with the cable category leading to a higher overall mIoU, as shown in Table 3.



Fig. 13. Qualitative results of DeepLabv3Plus (DLv3P)-ResNet101 on four categories of the TTPLA dataset. The models are evaluated at different stages: base (B), enhanced (E), and mobile (M).

6.5. Mobile model performance

In this subsection, we evaluate the performance of the DeepLabv3Plus models after the quantization and deployment on mobile devices. The quantitative results are presented in Table 3 comparing the models' performance, while Fig. 13 illustrating the qualitative results showcase segmentation outputs for the visual quality inspection task.

From Table 3, we observe that, when deployed on mobile devices, the mobile model variants experience a slight performance drop due to the quantization and optimization process, which reduces the model's size and computational complexity to make it suitable for mobile inference. The mobile DeepLabv3Plus-MobileNetv2 (DLv3P-MobileNetv2-M) model, specifically designed for mobile deployment, maintains a competitive mIoU of 75.48% while having significantly fewer parameters (3.51M) and a smaller model size (13.39 MB) compared to the ResNet-based models.

Qualitatively, Fig. 13 visually compares the segmentation results produced by the base, enhanced, and mobile versions of the DeepLabv3Plus model with the ResNet101 backbone. The first column shows the input images, while the second column demonstrates the corresponding ground truth annotations. The subsequent columns demonstrate the segmentation outputs of the base, enhanced, and mobile models, respectively. The qualitative results show that the enhanced DeepLabv3Plus-ResNet101 (DLv3P-ResNet101-E) model produces highly accurate segmentation masks, closely resembling the ground truth annotations. Especially, the enhanced model is less prone to perplexed objects, as shown in Fig. 13b. The segmentation quality of the mobile DeepLabv3Plus-ResNet101 (DLv3P-ResNet101-M) model is slightly lower compared to the enhanced version, with some minor discrepancies in the segmented regions. However, the mobile model still maintains good visual quality, successfully identifying and segmenting key objects of interest.

The results confirm the visual quality of the segmentation outputs produced by the mobile models, highlighting their effectiveness in real-world scenarios. Overall, our procedure of developing mobile models can provide a balance between performance and computational efficiency.

6.6. Textual explanation

Template 1: Prompt template for the GPT-4 Vision

System Message: You are an Explainable AI expert for semantic segmentation models.

The first image is the original image. The second image is the ground truth of category {category} in the original image. The third image is the prediction of an AI model for category {category} in the original image. The fourth image is the explanation map of the category {category} in the original image. Think step by step to understand how the explanation map and prediction align with the ground truth. First, capture the image context in the original image. Secondly, identify which parts belong to category {category} in the ground truth. Then, you look at the explanation map to see the saliency map for the segmentation mask of the category {category}. Your task is to check if the focused region in the explanation map supports the prediction for category {category}. Your final answer must be concise, simple, and separated by bullet points. First, briefly describe the explanation map. Secondly, describe the most focused region of category {category} in the explanation. Thirdly, describe the least focused region of category {category} in the explanation. Fourthly, assess the localization quality if the focused region aligns with the prediction and ground truth for category {category}.

User Message:

image_url: "data:image/jpeg;base64,{input_image}"
image_url: "data:image/jpeg;base64,{ground_truth_image}"
image_url: "data:image/jpeg;base64,{segmentation_image}"
image_url: "data:image/jpeg;base64,{explanation_image}"

In this subsection, we present the results of generating textual explanations to provide a human-readable interpretation of the segmentation results on the TTPLA dataset. The textual explanations aim to enhance the interpretability and trustworthiness of the mobile app for field engineers by offering a concise and intuitive summary of the model's rationale.

To generate the textual explanations, we design a prompt template that includes the original image, the ground truth image, the segmentation image, and the saliency map. The LVLM processes this information and leverages its pre-trained knowledge to generate human-readable explanations. Template 1 presents the prompt template used for generating these explanations. The system message informs LVLM that it should act as an XAI expert in describing saliency maps retrieved from XAI methods for a semantic segmentation model. It outlines the input structure, including the original image, the ground truth image, the segmentation image, and the explanation map image. The system message also instructs the model to think step-by-step, identifying parts of each image's specific category. Furthermore, it guides the model in describing the concentrated regions of the explanation map. It provides instructions on how to format the final answer, emphasizing correctness and simplicity for end-user understanding. The user message, on the other hand, provides the specific inputs for the model to process. In our framework, the user message includes the URLs of the original image, ground truth image, segmentation image, and explanation image. These images are passed to the model as base64-encoded strings.

Template 2 shows examples of the textual explanations generated by the LVLM for the good segmentation of *tower_wooden* (Fig. 14) and bad segmentation of *cable* (Fig. 15) categories, respectively. The generated textual explanation provides a clear overview of the saliency map, indicating areas of interest where the prediction model concentrates on identifying the presence of a wooden tower structure or cables. The explanation highlights the most concentrated region and points out the least concentrated region. Moreover, the textual explanation can also assess the model's segmentation performance, whether the model has a good or bad segmentation, by combining the information from the saliency image, prediction image, and ground truth image. The textual explanation shown directly on the mobile devices and the saliency map aims to support end-users in understanding the model's rationale in a human-centered manner.

7. Experiment 2: A data-centric approach

Based on the findings from Experiment 1, we further evaluate our framework in Experiment 2 using a data-centric approach. The experimental procedures remain consistent with those used in the previous experiment. However, this time we directly apply \mathcal{X} as the explaining technique without conducting comparative analyses of different XAI methods. Additionally, we utilize a different public industrial dataset, focusing on the inspection of substation equipment.

7.1. Dataset

The Substation Equipment dataset comprises 1660 images of electric substations collected using handheld cameras, Automated guided vehicle (AGV)-mounted cameras, and fixed-location cameras [119]. It features 15 categories of substation equipment, with a total of 50,705 annotated objects. These categories include *open_blade_disconnect_switch*, *breaker*, *closed_blade_disconnect_switch*, *open_tandem_disconnect_switch*, *fuse_disconnect_switch*, *porcelain_pin_insulator*, *closed_tandem_disconnect_switch*, *fuse_transformer*, *lightning_arrester*, *recloser*, *power_ transformer*, *current_transformer*, *glass_disc_insulator*, and *tripolar_ disconnect_switch*. Table 5 summarizes the object characteristics by class, detailing counts, area percentages, and dimension ranges (width and height) for analysis.

The dataset provides annotations in two formats: VOC-style polygonal JSON files and segmentation masks in Portable Network Graphic images. In our experiment, we use the VOC-style polygonal JSON files for semantic annotations. Fig. 16 presents samples from several categories, such as recloser, power transformer, porcelain pin insulator, and tripolar disconnect switch.

The primary application of this dataset is to develop automated inspection systems for substations. Training semantic segmentation models on this dataset allows researchers to improve the accuracy and reliability of these systems, resulting in more efficient and effective monitoring of substation equipment.

7.2. Base model performance

In this experiment, we demonstrate the efficiency of our framework when applied to the DeepLabv3Plus model with a ResNet101 backbone. The model is trained on the training set D_{train} for 1000 epochs, and the accuracy for each category is presented in Table 6.

The base DeepLabv3Plus-ResNet101 model achieves promising results on the validation set, with an overall mIoU of 73.45%. However, the model's performance varies across different categories. Some categories, such as *breaker*, *open_blade_disconnect_switch*, and *fuse_disconnect_ switch* achieve high IoU scores above 90%. In contrast, others, like *porcelain_pin_insulator* and *closed_blade_disconnect_switch*, have lower scores of around 30% and 49%, respectively. These performance variations underscored the challenges posed by the diverse object categories in the employed dataset (see Fig. 17).



Model Output:

- The explanation map is a visual representation showing the importance of different regions of the image for predicting the category tower lattice.

- The most focused region in the explanation map seems to be the central part of the tower, with yellow to red gradients, indicating a high degree of importance according to the model.

- The least focused region in the explanation map appears to be the bottom and peripheral parts of the tower, shown by the cooler colors such as blue, which suggest these areas have less influence on the model's decision.

- The segmentation accuracy seems high, where the focused region corresponds well with the AI model's prediction and the ground truth. This alignment indicates that the model focuses on relevant parts of the image to predict the category tower_lattice.



Fig. 15. Bad segmentation - Category: cable.

Model Output:

- The explanation map is a visual representation that indicates which areas of the image the AI model deemed relevant when determining where the category "cable" is present.

- The most focused region of the cable in the explanation map is evident along the linear features within the road, which correlate with the yellow lane markings and the road's edges.

- The least focused region in the explanation appears to be the areas on the pavement free from markings and the natural vegetation areas on the right side of the image.

- The focused region in the explanation map does not perfectly align with the ground truth for the cable category. The prediction appears to have false positives along the lane markings. It does not accurately distinguish between the road markings and actual cables, whereas the ground truth demarcates straight lines purported to be cables adjacent to the road on both sides. Therefore, there is a discrepancy between the focused regions in the explanation map and the true location of cables, as shown in the ground truth image.

Table 5

Summary of object characteristics by class in the Substation dataset, showing the number of objects, area percentages, and dimensions (width and height in pixels) with minimum, maximum, and average values.

Class	#Obj.	Min/Max/Avg area (%)	Min/Max/Avg width (px)	Min/Max/Avg height (px)
porcelain_pin_insulator	26 450	0.00/16.16/0.08	2/2244/57	4/2291/64
closed_blade_disconnect_switch	5 225	0.00/16.22/0.35	3/1107/97	9/2003/216
glass_disc_insulator	3 180	0.00/0.78/0.05	3/283/66	6/399/62
tripolar_disconnect_switch	2 348	0.00/4.9/0.23	3/839/98	5/754/124
recloser	2 330	0.00/25.44/1.52	4/1486/251	3/1355/230
current_transformer	2 1 2 8	0.00/2.68/0.25	4/1232/101	6/761/125
lightning_arrester	1 974	0.00/1.78/0.06	4/397/40	4/963/84
open_tandem_disconnect_switch	1 596	0.00/5.12/0.09	4/659/100	4/1114/108
muffle	1 354	0.00/9.26/0.09	5/280/41	4/956/165
breaker	980	0.00/15.42/1.08	7/1487/156	6/1468/298

7.3. Model improvement via XAI-guided data augmentation

After evaluating the segmentation performance of the base model on the validation set $\mathcal{D}_{\text{val}},$ we analyze the corresponding explanation

images generated by RISE, the recommended XAI method \mathcal{X} identified in Experiment 1. RISE analysis reveals systematic failure patterns in object detection and segmentation across specific equipment categories, particularly *recloser*, *glass_disc_insulator*, *porcelain_pin_insulator*,



Fig. 16. Samples from the Substation Equipment dataset represent the main objects of categories in the green masks (a) recloser, (b) power_transformer, (c) porcelain_pin_insulator, (d) tripolar_disconnect_switch.



Fig. 17. List of input images, ground truth, segmentation of the base DeepLabv3Plus-ResNet101-B model, and the corresponding explanations for four categories of the Substation equipment dataset.

Table 6

Accuracy comparison of three DeepLabv3Plus models with the ResNet101 backbone at different stages: base (B), enhanced (E), and mobile (M) in terms of average IoU (%) for each category and mIoU (%) on the Substation equipment validation set.

Model	В	E	М
breaker	92.93	94.51	92.51
closed_blade_disconnect_switch	49.14	50.26	48.02
closed_tandem_disconnect_switch	83.11	84.34	82.04
current_transformer	62.23	63.25	61.07
fuse_disconnect_switch	90.87	92.77	90.03
glass_disc_insulator	53.55	55.42	53.08
lightning_arrester	67.87	69.28	67.41
muffle	88.27	90.06	87.92
open_blade_disconnect_switch	93.34	95.18	93.06
open_tandem_disconnect_switch	83.85	85.54	83.51
porcelain_pin_insulator	30.11	32.03	30.05
potential_transformer	78.54	80.72	78.01
power_transformer	85.73	87.79	85.50
recloser	68.11	70.40	67.50
tripolar_disconnect_switch	83.22	85.24	82.79
mIoU	73.45	75.79	72.58

current_transformer, lightning_arrester and *closed_blade_disconnect_switches,* resulting in IoU scores below 70%. As shown in (Fig. 20), the explanation maps indicate insufficient model activations on target objects, especially under challenging imaging conditions such as (1) small spatial dimensions relative to the image size (as depicted in Table 5), (2) low contrast due to similar visual features with surrounding infrastructure, or (3) truncated or out-of-frame object instances at image boundaries.

19

To address this issue, we employ a sequence of image augmentation techniques [120] to enhance the training data, as shown in Fig. 18. The augmentation pipeline includes horizontal flipping, horizontal shifting, padding, Gaussian noise addition, perspective transformation, Contrast Limited Adaptive Histogram Equalization (CLAHE), sharpening, and brightness and contrast adjustments. These techniques aim to improve the model's robustness and generalization ability by simulating real-world variations in the training data.

After applying the image augmentation techniques to the training set D_{train} , we retrain the model and implement the mobile optimization algorithm (Algorithm 1) to obtain the enhanced model. Table 6 shows that the enhanced model $\tilde{\Theta}$ improves across all categories, with an overall mIoU of 75.79%, surpassing the base model's performance. Table 6 presents the mobile model's performance. Despite undergoing quantization and pruning, the mobile model θ maintains a competitive overall mIoU of 72.58%, experiencing only a slight decrease compared to the enhanced model. Its performance across individual categories remained consistent, with the largest drop being less than 3% for the *porcelain_pin_insulator* category.

7.4. Textual explanation

To generate textual explanations, we follow the prompting structure presented in Template 1. Specifically, Template 3 provides the textual explanations generated by the LVLM for two segmentation cases: successful segmentation of an object in the *power_transformer* category (Fig. 19) and unsuccessful segmentation of an object in the *recloser* category (Fig. 20). The results show that the LVLM effectively provided human-understandable textual explanations for both cases.



Fig. 18. Augmentation techniques on the training set of the Substation equipment dataset.

In the successful segmentation case, the LVLM highlights the central structures and components of the *power_transformer* as the primary focus in the explanation map, while the ground and less critical background structures received minimal attention. It also notes that these focused regions partially supported the prediction, indicating that the model utilized broader contextual cues for accurate segmentation.

In the failure case, the LVLM highlights the labeled components and connectors on the left as the primary focus in the explanation map rather than the *recloser* itself. It observes that these focused regions did not support the prediction for the *recloser*, indicating a significant issue with the model's ability to recognize and segment this category.

8. Discussion

In this section, we discuss the key findings of our results, highlighting both the advantages and limitations of the proposed framework. We also offer a variety of suggestions for future development.

8.1. XAI-guided data augmentation for model enhancement and deployment on the edge

Our experimental results highlight the effectiveness of the proposed XAI-integrated Visual Quality Inspection framework for industrial assets in both cloud and edge computing contexts. Quantitative and qualitative analyses demonstrate that integrating XAI techniques significantly enhances the performance and interpretability of semantic segmentation models deployed on mobile devices.

A key aspect of our framework is the application of XAI methods, such as RISE, to guide data augmentation and improve model performance. By leveraging insights from the saliency maps generated by RISE, domain experts can refine training data annotations, particularly in areas where the model struggled with segmentation. This process significantly improved segmentation accuracy, especially for challenging categories. The enhanced DeepLabv3Plus-ResNet101 model's higher IoU scores compared to the base model highlight the value of integrating XAI into the model development process for edge deployment.

Another crucial feature of our framework is optimizing models for edge devices to enable real-time field inspections. Our results indicate that the quantized and pruned edge models maintain competitive performance while significantly reducing model size and computational requirements. This optimization is essential for efficient on-device inference, making the visual quality inspection system more practical and accessible for field engineers using edge devices with limited resources.

8.2. Human-centered textual explanation

Integrating textual explanations generated by LVLM greatly enhances the interpretability and applicability of edge-deployed applications. These clear and intuitive summaries of the model's decisionmaking process help field engineers better understand and trust the segmentation results on their mobile devices. Specifically, the explanations accurately highlight the most and least focused regions in the saliency maps, allowing users to effectively assess the quality of the model's predictions by comparing the explanation maps with ground truth annotations. Delivering interpretable insights directly on mobile devices is essential for field engineers to make informed decisions and confidently rely on the system's outputs.

Overall, our proposed framework demonstrates the feasibility and benefits of generating human-readable explanations alongside saliency maps at the edge, resulting in more transparent and trustworthy visual quality inspection systems.

8.3. Limitations and future developments

While our XAI-integrated Visual Quality Inspection framework shows promising results, there exist some limitations that should be considered.

Firstly, the XAI-guided annotation augmentation process relies on the expertise and manual effort of domain experts. Automating or semi-automating this process could enhance scalability and efficiency, particularly for large-scale edge deployments. Future work could explore techniques to automate the annotation refinement process based on XAI insights.

Secondly, generating textual explanations on edge devices may introduce latency and computational overhead. Although the current framework shows this approach is feasible, it also highlights the need for further research to address these drawbacks. This could involve exploring more lightweight language models or developing techniques specifically designed for edge computing scenarios.

Lastly, the framework's applicability to other visual inspection domains in edge computing scenarios remains to be explored. We recommend future work investigate the generalizability of the proposed approach across different visual quality inspection tasks and evaluate its performance and interpretability on various datasets and edge computing environments.

9. Conclusion

This paper introduces an XAI-integrated Visual Quality Inspection framework for industrial assets, meeting the increasing demand for explainable and trustworthy AI systems in edge computing environments. Our framework combines cutting-edge semantic segmentation models, XAI techniques, quantization algorithms, LVLM, and edge deployment to enhance both the performance and interpretability of visual quality inspection systems for end-users. Our experimental results demonstrated the remarkable effectiveness of XAI methods in data improvement and augmentation, which led to enhanced model accuracy and performance, especially for challenging and complex categories.



Model Output:

- The explanation map highlights several regions of importance, with some focus around the power transformer but also attention on other structures.

- The central structures and parts of the power transformer have significant focus in the explanation map.

- The ground and less critical background structures have the least focus.

- The focused regions in the explanation map partially support the prediction for the power transformer. The prediction is accurate, but the explanation map indicates that the model's focus is distributed, suggesting it uses broader contextual cues.



Fig. 20. Bad segmentation - Category: recloser.

Model Output:

- The explanation map highlights regions of importance, with significant focus on the left-side components and connectors, rather than the recloser.

- The labeled components and connectors on the left are the most focused regions in the explanation map.

- The background and some parts of the ground and structure are the least focused regions.

- The focused regions in the explanation map do not support the prediction for the recloser. The model failed to segment the recloser, indicating a significant issue with the model's ability to recognize and segment this category.

A notable feature of our framework is the enhancement module, which optimizes the model for edge devices through quantization and pruning, making it more lightweight and efficient while maintaining competitive performance. Additionally, integrating textual explanations generated by LVLM significantly enhances the interpretability of the edge-deployed application. Last but not least, providing interpretable insights directly on edge devices empowers field engineers to make informed decisions and increases their confidence in the system's outputs.

On the other hand, the framework has certain limitations, including the human-in-the-loop effort required for XAI-guided annotation augmentation and the potential latency and computational overhead introduced by generating explanations on edge devices. Future work could explore methods to automate the annotation refinement process and optimize the explanation generation stage for edge computing scenarios.

In conclusion, despite all the limitations discussed, our work has successfully paved the way for more transparent, trustworthy, and effective industrial visual inspection systems deployed at the edge. By addressing key challenges and demonstrating the feasibility of our approach, we lay a solid foundation for future advancements. This can significantly enhance the reliability and performance of edge-based visual inspection technologies, ultimately benefiting both end-users and developers by providing more accurate, interpretable, and efficient solutions.

CRediT authorship contribution statement

Hung Truong Thanh Nguyen: Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Resources, Project administration, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. Loc Phuc Truong Nguyen: Writing – original draft, Validation, Resources, Formal analysis, Investigation, Methodology, Software. Hung Cao: Writing – review & editing, Writing – original draft, Supervision, Resources, Project administration, Funding acquisition, Conceptualization, Methodology.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This work was funded by the NBIF Talent Recruitment Fund (TRF2003-001) and the UNB-FCS Startup Fund (22–23 START UP/

H CAO). The computing resource used for this study was partially supported by CFI Project Number 39473 - Smart Campus Integration and Testing (SCIT Lab).

Data availability

The datasets used in this research are public and specified in the manuscript.

References

- [1] M. Garouani, A. Ahmad, M. Bouneffa, M. Hamlich, G. Bourguin, A. Lewandowski, Towards big industrial data mining through explainable automated machine learning, Int. J. Adv. Manuf. Technol. 120 (1) (2022) 1169–1188.
- [2] T.T.H. Nguyen, V.B. Truong, V.T.K. Nguyen, Q.H. Cao, Q.K. Nguyen, Towards trust of explainable ai in thyroid nodule diagnosis, in: International Workshop on Health Intelligence, Springer, 2023, pp. 11–26.
- [3] T.-Y. Wu, Y.-T. Wang, Locally interpretable one-class anomaly detection for credit card fraud detection, in: 2021 International Conference on Technologies and Applications of Artificial Intelligence, TAAI, IEEE, 2021, pp. 25–30.
- [4] 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.
- [5] V. Bento, M. Kohler, P. Diaz, L. Mendoza, M.A. Pacheco, Improving deep learning performance by using explainable artificial intelligence (XAI) approaches, Discov. Artif. Intell. 1 (1) (2021) 9, http://dx.doi.org/10.1007/s44163-021-00008-y, URL https://doi.org/10.1007/s44163-021-00008-y.
- [6] S. Teso, K. Kersting, Explanatory interactive machine learning, in: Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society, AIES '19, Association for Computing Machinery, New York, NY, USA, 2019, pp. 239–245, http://dx.doi.org/10.1145/3306618.3314293.
- [7] S.A. Bargal, A. Zunino, V. Petsiuk, J. Zhang, K. Saenko, V. Murino, S. Sclaroff, Guided zoom: Questioning network evidence for fine-grained classification, 2020, arXiv:1812.02626.
- [8] L. Weber, S. Lapuschkin, A. Binder, W. Samek, Beyond explaining: Opportunities and challenges of XAI-based model improvement, Inf. Fusion (2022).
- [9] T. Clement, H.T.T. Nguyen, N. Kemmerzell, M. Abdelaal, D. Stjelja, Coping with data distribution shifts: XAI-based adaptive learning with SHAP clustering for energy consumption prediction, in: Australasian Joint Conference on Artificial Intelligence, Springer, 2023, pp. 147–159.
- [10] S.-K. Yeom, P. Seegerer, S. Lapuschkin, A. Binder, S. Wiedemann, K.-R. Müller, W. Samek, Pruning by explaining: A novel criterion for deep neural network pruning, Pattern Recognit. 115 (2021) 107899, http://dx.doi.org/10.1016/ j.patcog.2021.107899, URL https://www.sciencedirect.com/science/article/pii/ S0031320321000868.
- [11] M. Sabih, F. Hannig, J. Teich, Utilizing explainable AI for quantization and pruning of deep neural networks, 2020, arXiv:2008.09072.
- [12] D. Becking, M. Dreyer, W. Samek, K. Müller, S. Lapuschkin, ECQ: Explainabilitydriven quantization for low-bit and sparse DNNs, in: A. Holzinger, R. Goebel, R. Fong, T. Moon, K.-R. Müller, W. Samek (Eds.), XxAI - beyond Explainable AI: International Workshop, Held in Conjunction with ICML 2020, July 18, 2020, Vienna, Austria, Revised and Extended Papers, Springer International Publishing, Cham, 2022, pp. 271–296, http://dx.doi.org/10.1007/978-3-031-04083-2_14.
- [13] J. Gama, S. Nowaczyk, S. Pashami, R.P. Ribeiro, G.J. Nalepa, B. Veloso, XAI for predictive maintenance, in: Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, KDD '23, Association for Computing Machinery, New York, NY, USA, 2023, pp. 5798–5799, http: //dx.doi.org/10.1145/3580305.3599578.
- [14] C. Molnar, Interpretable Machine Learning, second ed., 2022, URL https:// christophm.github.io/interpretable-ml-book.
- [15] A.B. Arrieta, N. Díaz-Rodríguez, J. Del Ser, A. Bennetot, S. Tabik, A. Barbado, S. García, S. Gil-López, D. Molina, R. Benjamins, et al., Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI, Inf. Fusion 58 (2020) 82–115.
- [16] D. Zha, Z.P. Bhat, K.-H. Lai, F. Yang, X. Hu, Data-centric ai: Perspectives and challenges, in: Proceedings of the 2023 SIAM International Conference on Data Mining, SDM, SIAM, 2023, pp. 945–948.
- [17] F. Bodria, F. Giannotti, R. Guidotti, F. Naretto, D. Pedreschi, S. Rinzivillo, Benchmarking and survey of explanation methods for black box models, Data Min. Knowl. Discov. 37 (5) (2023) 1719–1778.
- [18] M.P.S. Lorente, E.M. Lopez, L.A. Florez, A.L. Espino, J.A.I. Martínez, A.S. de Miguel, Explaining deep learning-based driver models, Appl. Sci. 11 (8) (2021) 3321.

- [19] Z. Li, Leveraging explainable artificial intelligence and big trip data to understand factors influencing willingness to ridesharing, Travel Behav. Soc. 31 (2023) 284–294, http://dx.doi.org/10.1016/j.tbs.2022.12.006, URL https: //www.sciencedirect.com/science/article/pii/S2214367X22001466.
- [20] M. Dikmen, C. Burns, The effects of domain knowledge on trust in explainable AI and task performance: A case of peer-to-peer lending, Int. J. Hum.-Comput. Stud. 162 (2022) 102792.
- [21] X. Dastile, T. Celik, Making deep learning-based predictions for credit scoring explainable, IEEE Access 9 (2021) 50426–50440, http://dx.doi.org/10.1109/ ACCESS.2021.3068854.
- [22] B. Žlahtič, J. Završnik, H. Blažun Vošner, P. Kokol, D. Šuran, T. Završnik, Agile machine learning model development using data canyons in medicine: A step towards explainable artificial intelligence and flexible expert-based model improvement, Appl. Sci. 13 (14) (2023) 8329.
- [23] P. Guleria, P. Naga Srinivasu, S. Ahmed, N. Almusallam, F.K. Alarfaj, XAI framework for cardiovascular disease prediction using classification techniques, Electronics 11 (24) (2022) http://dx.doi.org/10.3390/electronics11244086, URL https://www.mdpi.com/2079-9292/11/24/4086.
- [24] E. Melo, I. Silva, D.G. Costa, C.M. Viegas, T.M. Barros, On the use of explainable artificial intelligence to evaluate school dropout, Educ. Sci. 12 (12) (2022) 845.
- [25] N. Nur, A. Benedict, O. Eltayeby, W. Dou, M. Dorodchi, X. Niu, M. Maher, C. Chambers, Explainable ai for data driven learning analytics: A holistic approach to engage advisors in knowledge discovery, in: EDULEARN22 Proceedings, IATED, 2022, pp. 10300–10306.
- [26] M. Tsiakmaki, O. Ragos, A case study of interpretable counterfactual explanations for the task of predicting student academic performance, in: 2021 25th International Conference on Circuits, Systems, Communications and Computers, CSCC, IEEE, 2021, pp. 120–125.
- [27] Y. Kardovskyi, S. Moon, Artificial intelligence quality inspection of steel bars installation by integrating mask R-CNN and stereo vision, Autom. Constr. 130 (2021) 103850.
- [28] M.T.G. Diaz, D. Ghosh, A. Arantes, H. Khorosgani, M. Alam, G. Sin, C. Gupta, Guided visual inspection enabled by AI-based detection models, in: 2021 IEEE International Conference on Prognostics and Health Management, ICPHM, IEEE, 2021, pp. 1–8.
- [29] C. Eiras-Franco, B. Guijarro-Berdiñas, A. Alonso-Betanzos, A. Bahamonde, A scalable decision-tree-based method to explain interactions in dyadic data, Decis. Support Syst. 127 (2019) 113141, http://dx.doi.org/10.1016/ j.dss.2019.113141, URL https://www.sciencedirect.com/science/article/pii/ S0167923619301708.
- [30] M.A. Islam, D.T. Anderson, A.J. Pinar, T.C. Havens, G. Scott, J.M. Keller, Enabling explainable fusion in deep learning with fuzzy integral neural networks, IEEE Trans. Fuzzy Syst. 28 (7) (2020) 1291–1300, http://dx.doi.org/10.1109/ TFUZZ.2019.2917124.
- [31] Á.L. de Santana, C.R. Francês, C.A. Rocha, S.V. Carvalho, N.L. Vijaykumar, L.P. Rego, J.C. Costa, Strategies for improving the modeling and interpretability of Bayesian networks, Data Knowl. Eng. 63 (1) (2007) 91–107, http:// dx.doi.org/10.1016/j.datak.2006.10.005, URL https://www.sciencedirect.com/ science/article/pii/S0169023X06001959. Data Warehouse and Knowledge Discovery (DAWAK '05).
- [32] Y. Rong, T. Leemann, T. Nguyen, L. Fiedler, P. Qian, V. Unhelkar, T. Seidel, G. Kasneci, E. Kasneci, Towards human-centered explainable AI: A survey of user studies for model explanations, IEEE Trans. Pattern Anal. Mach. Intell. 46 (04) (2024) 2104–2122, http://dx.doi.org/10.1109/TPAMI.2023.3331846.
- [33] T. Clement, N. Kemmerzell, M. Abdelaal, M. Amberg, XAIR: A systematic metareview of explainable AI (XAI) aligned to the software development process, Mach. Learn. Knowl. Extract. 5 (1) (2023) 78–108, URL https://www. mdpi.com/2504-4990/5/1/6.
- [34] L. Weber, S. Lapuschkin, A. Binder, W. Samek, Beyond explaining: Opportunities and challenges of XAI-based model improvement, Inf. Fusion 92 (2023) 154–176, http://dx.doi.org/10.1016/j.inffus.2022.11.013, URL https:// www.sciencedirect.com/science/article/pii/S1566253522002238.
- [35] J. Sun, S. Lapuschkin, W. Samek, Y. Zhao, N. Cheung, A. Binder, Explanationguided training for cross-domain few-shot classification, in: 2020 25th International Conference on Pattern Recognition, ICPR, IEEE Computer Society, Los Alamitos, CA, USA, 2021, pp. 7609–7616, http://dx.doi.org/10.1109/ ICPR48806.2021.9412941, URL https://doi.ieeecomputersociety.org/10.1109/ ICPR48806.2021.9412941.
- [36] A. Zunino, S. Bargal, R. Volpi, M. Sameki, J. Zhang, S. Sclaroff, V. Murino, K. Saenko, Explainable deep classification models for domain generalization, in: 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops, CVPRW, IEEE Computer Society, Los Alamitos, CA, USA, 2021, pp. 3227–3236, http://dx.doi.org/10.1109/CVPRW53098.2021.00361, URL https: //doi.ieeecomputersociety.org/10.1109/CVPRW53098.2021.00361.
- [37] J. Blunk, N. Penzel, P. Bodesheim, J. Denzler, Beyond debiasing: Actively steering feature selection via loss regularization, in: U. Köthe, C. Rother (Eds.), Pattern Recognition, Springer Nature Switzerland, Cham, 2024, pp. 394–408.

- [38] F. Liu, B. Avci, Incorporating priors with feature attribution on text classification, in: A. Korhonen, D. Traum, L. Màrquez (Eds.), Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, Association for Computational Linguistics, Florence, Italy, 2019, pp. 6274–6283, http://dx. doi.org/10.18653/v1/P19-1631, URL https://aclanthology.org/P19-1631.
- [39] L. Rieger, C. Singh, W.J. Murdoch, B. Yu, Interpretations are useful: penalizing explanations to align neural networks with prior knowledge, in: Proceedings of the 37th International Conference on Machine Learning, ICML '20, JMLR.org, 2020.
- [40] J. Zhang, S.A. Bargal, Z. Lin, J. Brandt, X. Shen, S. Sclaroff, Top-down neural attention by excitation backprop, Int. J. Comput. Vis. 126 (10) (2018) 1084–1102, http://dx.doi.org/10.1007/s11263-017-1059-x.
- [41] G. Erion, J.D. Janizek, P. Sturmfels, S.M. Lundberg, S.-I. Lee, Improving performance of deep learning models with axiomatic attribution priors and expected gradients, Nat. Mach. Intell. 3 (7) (2021) 620–631, http://dx.doi.org/ 10.1038/s42256-021-00343-w.
- [42] V. Nagisetty, L. Graves, J. Scott, V. Ganesh, xAI-GAN: Enhancing generative adversarial networks via explainable ai systems, 2022, arXiv:2002.10438.
- [43] K. Simonyan, A. Zisserman, Very deep convolutional networks for large-scale image recognition, 2014, arXiv preprint arXiv:1409.1556.
- [44] J. Redmon, S. Divvala, R. Girshick, A. Farhadi, You only look once: Unified, real-time object detection, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2016, pp. 779–788.
- [45] K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2016, pp. 770–778.
- [46] A.G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, H. Adam, Mobilenets: Efficient convolutional neural networks for mobile vision applications, 2017, arXiv preprint arXiv:1704.04861.
- [47] J. Long, E. Shelhamer, T. Darrell, Fully convolutional networks for semantic segmentation, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2015, pp. 3431–3440.
- [48] A. Howard, M. Sandler, G. Chu, L.-C. Chen, B. Chen, M. Tan, W. Wang, Y. Zhu, R. Pang, V. Vasudevan, et al., Searching for mobilenetv3, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, 2019, pp. 1314–1324.
- [49] L.-C. Chen, G. Papandreou, F. Schroff, H. Adam, Rethinking atrous convolution for semantic image segmentation, 2017, arXiv:1706.05587.
- [50] 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, 2016, arXiv 2016. arXiv preprint arXiv:1606.00915. 1.
- [51] Z. Yang, Semantic segmentation method based on improved DeeplabV3+, in: International Conference on Cloud Computing, Performance Computing, and Deep Learning, CCPCDL 2023, Vol. 12712, SPIE, 2023, pp. 32–37.
- [52] 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, Voll 34, 2020, pp. 13943–13944.
- [53] S.-A. Rebuffi, R. Fong, X. Ji, A. Vedaldi, There and back again: Revisiting backpropagation saliency methods, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2020, pp. 8839–8848.
- [54] S. Bach, A. Binder, G. Montavon, F. Klauschen, K.-R. Müller, W. Samek, On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation, PLoS One 10 (7) (2015) e0130140.
- [55] A. Shrikumar, P. Greenside, A. Kundaje, Learning important features through propagating activation differences, in: International Conference on Machine Learning, PMLR, 2017, pp. 3145–3153.
- [56] 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.
- [57] A. Chattopadhay, A. Sarkar, P. Howlader, V.N. Balasubramanian, Grad-cam++: Generalized gradient-based visual explanations for deep convolutional networks, in: 2018 IEEE Winter Conference on Applications of Computer Vision, WACV, IEEE, 2018, pp. 839–847.
- [58] R.R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, D. Batra, Gradcam: Visual explanations from deep networks via gradient-based localization, in: Proceedings of the IEEE International Conference on Computer Vision, 2017, pp. 618–626.
- [59] H. Wang, Z. Wang, M. Du, F. Yang, Z. Zhang, S. Ding, P. Mardziel, X. Hu, Score-CAM: Score-weighted visual explanations for convolutional neural networks, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops, 2020, pp. 24–25.
- [60] H.G. Ramaswamy, et al., 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.
- [61] M.B. Muhammad, M. Yeasin, Eigen-cam: Class activation map using principal components, in: 2020 International Joint Conference on Neural Networks, IJCNN, IEEE, 2020, pp. 1–7.
- [62] R. Fu, Q. Hu, X. Dong, Y. Guo, Y. Gao, B. Li, Axiom-based grad-cam: Towards accurate visualization and explanation of cnns, 2020, arXiv preprint arXiv: 2008.02312.

- [63] P.X. Nguyen, H.Q. Cao, K.V. Nguyen, H. Nguyen, T. Yairi, SeCAM: Tightly accelerate the image explanation via region-based segmentation, IEICE Trans. Inf. Syst. 105 (8) (2022) 1401–1417.
- [64] 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.
- [65] Q.K. Nguyen, T.T.H. Nguyen, V.T.K. Nguyen, V.B. Truong, T. Phan, H. Cao, Efficient and concise explanations for object detection with Gaussian-class activation mapping explainer, 2024, arXiv preprint arXiv:2404.13417.
- [66] M.D. Zeiler, R. Fergus, Visualizing and understanding convolutional networks, in: Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part I 13, Springer, 2014, pp. 818–833.
- [67] M.T. Ribeiro, S. Singh, C. Guestrin, "Why should I trust you?": Explaining the predictions of any classifier, 2016, arXiv:1602.04938.
- [68] V. Petsiuk, A. Das, K. Saenko, Rise: Randomized input sampling for explanation of black-box models, 2018, arXiv preprint arXiv:1806.07421.
- [69] V. Petsiuk, R. Jain, V. Manjunatha, V.I. Morariu, A. Mehra, V. Ordonez, K. Saenko, Black-box explanation of object detectors via saliency maps, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2021, pp. 11443–11452.
- [70] Q. Yang, X. Zhu, J.-K. Fwu, Y. Ye, G. You, Y. Zhu, Mfpp: Morphological fragmental perturbation pyramid for black-box model explanations, in: 2020 25th International Conference on Pattern Recognition, ICPR, IEEE, 2021, pp. 1376–1383.
- [71] V.B. Truong, T.T.H. Nguyen, V.T.K. Nguyen, Q.K. Nguyen, Q.H. Cao, Towards better explanations for object detection, in: B. Yanıkoğlu, W. Buntine (Eds.), Proceedings of the 15th Asian Conference on Machine Learning, in: Proceedings of Machine Learning Research, vol. 222, PMLR, 2024, pp. 1385–1400, URL https://proceedings.mlr.press/v222/truong24a.html.
- [72] 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.
- [73] N.F. Heide, E. Müller, J. Petereit, M. Heizmann, X 3 SEG: model-agnostic explanations for the semantic segmentation of 3D point clouds with prototypes and criticism, in: 2021 IEEE International Conference on Image Processing, ICIP, IEEE, 2021, pp. 3687–3691.
- [74] R.L. Draelos, L. Carin, Hirescam: Faithful location representation in visual attention for explainable 3d medical image classification, 2020, arXiv preprint arXiv:2011.08891.
- [75] J. Gildenblat, contributors, Pytorch library for CAM methods, 2021, https: //github.com/jacobgil/pytorch-grad-cam.
- [76] H. Tang, Manufacturing System and Process Development for Vehicle Assembly, SAE International, 2017.
- [77] X. Sun, J. Gu, S. Tang, J. Li, Research progress of visual inspection technology of steel products—a review, Appl. Sci. 8 (11) (2018) 2195.
- [78] A.Q. Md, K. Jha, S. Haneef, A.K. Sivaraman, K.F. Tee, A review on data-driven quality prediction in the production process with machine learning for industry 4.0, Processes 10 (10) (2022) 1966.
- [79] Y.D. Yasuda, F.A. Cappabianco, L.E.G. Martins, J.A. Gripp, Aircraft visual inspection: A systematic literature review, Comput. Ind. 141 (2022) 103695.
- [80] M. Ilchuk, A. Stadnyk, Computer visual inspection of pear quality, Meas. Equip. Metrol. (2023) http://dx.doi.org/10.23939/istcmtm2023.01.025.
- [81] J.M. Rožanec, B. Šircelj, B. Fortuna, D. Mladenić, Adaptive explainable artificial intelligence for visual defect inspection, Procedia Comput. Sci. 232 (2024) 3034–3043.
- [82] F. Lupi, M. Biancalana, A. Rossi, M. Lanzetta, A framework for flexible and reconfigurable vision inspection systems, Int. J. Adv. Manuf. Technol. 129 (1–2) (2023) 871–897.
- [83] H. Gunraj, P. Guerrier, S. Fernandez, A. Wong, SolderNet: Towards trustworthy visual inspection of solder joints in electronics manufacturing using explainable artificial intelligence, in: Proceedings of the AAAI Conference on Artificial Intelligence, Vol. 37, 2023, pp. 15668–15674.
- [84] R. Hoffmann, C. Reich, A systematic literature review on artificial intelligence and explainable artificial intelligence for visual quality assurance in manufacturing, Electronics 12 (22) (2023) 4572.
- [85] İ. Kök, Y. Ergun, N. Uğur, Explainable AI-powered edge computing solution for smart building energy management in green IoT, in: Low-Cost Digital Solutions for Industrial Automation, LoDiSA 2023, Vol. 2023, IET, 2023, pp. 150–157.
- [86] S. Garg, K. Kaur, G.S. Aujla, G. Kaddoum, P. Garigipati, M. Guizani, Trusted explainable AI for 6G-enabled edge cloud ecosystem, IEEE Wirel. Commun. 30 (3) (2023) 163–170.
- [87] J. Dutta, D. Puthal, Human-centered explainable ai at the edge for ehealth, in: 2023 IEEE International Conference on Edge Computing and Communications, EDGE, IEEE, 2023, pp. 227–232.
- [88] D. Wang, Q. Yang, A. Abdul, B.Y. Lim, Designing theory-driven user-centric explainable AI, in: Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, 2019, pp. 1–15.

- [89] S. Yu, Towards trustworthy and understandable AI: Unraveling explainability strategies on simplifying algorithms, appropriate information disclosure, and high-level collaboration, in: Proceedings of the 26th International Academic Mindtrek Conference, 2023, pp. 133–143.
- [90] A. Bertrand, T. Viard, R. Belloum, J.R. Eagan, W. Maxwell, On selective, mutable and dialogic XAI: A review of what users say about different types of interactive explanations, in: Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems, 2023, pp. 1–21.
- [91] J.-P. Poli, W. Ouerdane, R. Pierrard, Generation of textual explanations in XAI: The case of semantic annotation, in: 2021 IEEE International Conference on Fuzzy Systems, FUZZ-IEEE, IEEE, 2021, pp. 1–6.
- [92] D.H. Park, L.A. Hendricks, Z. Akata, A. Rohrbach, B. Schiele, T. Darrell, M. Rohrbach, Multimodal explanations: Justifying decisions and pointing to the evidence, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2018, pp. 8779–8788.
- [93] L.A. Hendricks, R. Hu, T. Darrell, Z. Akata, Grounding visual explanations, in: Proceedings of the European Conference on Computer Vision, ECCV, 2018, pp. 264–279.
- [94] K. Xu, J. Ba, R. Kiros, K. Cho, A. Courville, R. Salakhudinov, R. Zemel, Y. Bengio, Show, attend and tell: Neural image caption generation with visual attention, in: International Conference on Machine Learning, PMLR, 2015, pp. 2048–2057.
- [95] J. Kim, A. Rohrbach, T. Darrell, J. Canny, Z. Akata, Textual explanations for self-driving vehicles, in: Proceedings of the European Conference on Computer Vision, ECCV, 2018, pp. 563–578.
- [96] W. Dai, J. Li, D. Li, A.M.H. Tiong, J. Zhao, W. Wang, B. Li, P.N. Fung, S. Hoi, Instructblip: Towards general-purpose vision-language models with instruction tuning, Adv. Neural Inf. Process. Syst. 36 (2024).
- [97] T. Brown, B. Mann, N. Ryder, M. Subbiah, J.D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, et al., Language models are few-shot learners, Adv. Neural Inf. Process. Syst. 33 (2020) 1877–1901.
- [98] A. Chowdhery, S. Narang, J. Devlin, M. Bosma, G. Mishra, A. Roberts, P. Barham, H.W. Chung, C. Sutton, S. Gehrmann, et al., Palm: Scaling language modeling with pathways, J. Mach. Learn. Res. 24 (240) (2023) 1–113.
- [99] Z. Peng, W. Wang, L. Dong, Y. Hao, S. Huang, S. Ma, F. Wei, Kosmos-2: Grounding multimodal large language models to the world, 2023, arXiv preprint arXiv:2306.14824.
- [100] A. Awadalla, I. Gao, J. Gardner, J. Hessel, Y. Hanafy, W. Zhu, K. Marathe, Y. Bitton, S. Gadre, S. Sagawa, et al., Openflamingo: An open-source framework for training large autoregressive vision-language models, 2023, arXiv preprint arXiv:2308.01390.
- [101] R. Bavishi, E. Elsen, C. Hawthorne, M. Nye, A. Odena, A. Somani, S. Taşırlar, Introducing our multimodal models, 2023, URL https://www.adept.ai/blog/ fuyu-8b.
- [102] X. Chen, X. Wang, S. Changpinyo, A. Piergiovanni, P. Padlewski, D. Salz, S. Goodman, A. Grycner, B. Mustafa, L. Beyer, et al., Pali: A jointly-scaled multilingual language-image model, 2022, arXiv preprint arXiv:2209.06794.
- [103] OpenAI, GPT-4V(ision) system card, 2023, https://cdn.openai.com/papers/ GPTV_System_Card.pdf. (Accessed 14 February 2024).
- [104] X. Dong, P. Zhang, Y. Zang, Y. Cao, B. Wang, L. Ouyang, S. Zhang, H. Duan, W. Zhang, Y. Li, et al., InternLM-XComposer2-4KHD: A pioneering large visionlanguage model handling resolutions from 336 pixels to 4K HD, 2024, arXiv preprint arXiv:2404.06512.

- [105] D. Zhu, J. Chen, X. Shen, X. Li, M. Elhoseiny, Minigpt-4: Enhancing visionlanguage understanding with advanced large language models, 2023, arXiv preprint arXiv:2304.10592.
- [106] H. Nguyen, T. Clement, L. Nguyen, N. Kemmerzell, B. Truong, K. Nguyen, M. Abdelaal, H. Cao, LangXAI: Integrating large vision models for generating textual explanations to enhance explainability in visual perception tasks, in: K. Larson (Ed.), Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, IJCAI-24, International Joint Conferences on Artificial Intelligence Organization, 2024, pp. 8754–8758, http://dx.doi.org/10.24963/ ijcai.2024/1025.
- [107] Y. Chen, X-iqe: explainable image quality evaluation for text-to-image generation with visual large language models, 2023, arXiv preprint arXiv:2305. 10843.
- [108] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, L. Fei-Fei, Imagenet: A large-scale hierarchical image database, in: 2009 IEEE Conference on Computer Vision and Pattern Recognition, IEEE, 2009, pp. 248–255.
- [109] C.H. Sudre, W. Li, T. Vercauteren, S. Ourselin, M. Jorge Cardoso, Generalised dice overlap as a deep learning loss function for highly unbalanced segmentations, in: Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support: Third International Workshop, DLMIA 2017, and 7th International Workshop, ML-CDS 2017, Held in Conjunction with MICCAI 2017, Québec City, QC, Canada, September 14, Proceedings 3, Springer, 2017, pp. 240–248.
- [110] A.H. Murphy, The finley affair: A signal event in the history of forecast verification, Weather Forecast. 11 (1) (1996) 3–20.
- [111] D.P. Kingma, Adam: A method for stochastic optimization, 2014, arXiv preprint arXiv:1412.6980.
- [112] K. Schulz, L. Sixt, F. Tombari, T. Landgraf, Restricting the flow: Information bottlenecks for attribution, 2020, arXiv preprint arXiv:2001.00396.
- [113] W. Samek, A. Binder, G. Montavon, S. Lapuschkin, K.-R. Müller, Evaluating the visualization of what a deep neural network has learned, IEEE Trans. Neural Netw. Learn. Syst. 28 (11) (2016) 2660–2673.
- [114] S. Hooker, D. Erhan, P.-J. Kindermans, B. Kim, A benchmark for interpretability methods in deep neural networks, Adv. Neural Inf. Process. Syst. 32 (2019).
- [115] V. Bento, M. Kohler, P. Diaz, L. Mendoza, M.A. Pacheco, Improving deep learning performance by using explainable artificial intelligence (XAI) approaches, Discov. Artif. Intell. 1 (2021) 1–11.
- [116] Z. Yang, L. Li, K. Lin, J. Wang, C.-C. Lin, Z. Liu, L. Wang, The dawn of LMMs: Preliminary explorations with GPT-4V(ision), 2023, arXiv:2309.17421.
- [117] X. Zhang, Y. Lu, W. Wang, A. Yan, J. Yan, L. Qin, H. Wang, X. Yan, W.Y. Wang, L.R. Petzold, GPT-4V(ision) as a generalist evaluator for vision-language tasks, 2023, arXiv:2311.01361.
- [118] R. Abdelfattah, X. Wang, S. Wang, TTPLA: An aerial-image dataset for detection and segmentation of transmission towers and power lines, in: Proceedings of the Asian Conference on Computer Vision, 2020.
- [119] A. Gomes, A semantically annotated 15-class ground truth dataset for substation equipment, 2023, http://dx.doi.org/10.5281/zenodo.7884270.
- [120] A. Buslaev, V.I. Iglovikov, E. Khvedchenya, A. Parinov, M. Druzhinin, A.A. Kalinin, Albumentations: fast and flexible image augmentations, Information 11 (2) (2020) 125.