

Promptable Vision Foundation Models for Industrial Defect Segmentation and Quality Inspection

Dean Ortiz

Department of Computer Science, George Mason University, Fairfax, VA, USA.
deanortiz87@gmu.edu

Neeraj Mahajan

Department of Computer Science, Colorado State University, Fort Collins, CO, USA.
neeraj1987@colostate.edu

Mahesh Reddy

Department of Computer Science, University of Alabama at Birmingham, Birmingham, AL, USA.
reddy747@uab.edu

Abstract

Industrial defect segmentation and quality inspection are critical to maintaining product integrity across manufacturing sectors such as automotive, electronics, and aerospace. Traditional inspection systems rely on manually engineered features or supervised deep learning pipelines that require large, task-specific annotated datasets and frequent retraining when production conditions change. Promptable vision foundation models, such as the Segment Anything Model and vision-language architectures, offer a paradigm shift by enabling flexible, few-shot defect segmentation through natural language or visual prompts. This paper provides a system-level analysis of deploying such models in industrial environments, focusing on architectural trade-offs between generality and specificity, infrastructure requirements for real-time inference, and governance frameworks for model validation and fairness. We examine the interplay between model scale, latency, and on-premise versus cloud deployment, highlighting the sustainability implications of large-scale transformer architectures. Robustness to domain shifts, class imbalance, and adversarial inputs is discussed, alongside policy considerations for accountability and certification of AI-based inspection systems. Cross-domain comparisons with traditional machine vision and deep learning methods illustrate the conditions under which promptable models offer superior adaptability. The paper concludes by outlining future research directions, including lightweight model distillation, continuous learning under concept drift, and regulatory alignment for safety-critical quality assurance.

Keywords

promptable vision foundation models, industrial defect segmentation, quality inspection, few-shot learning, model deployment, robustness, governance, sustainability.

1. Introduction

The advent of foundation models in computer vision has fundamentally altered the landscape of visual understanding tasks. These models, pre-trained on vast and diverse datasets, exhibit remarkable generalization capabilities that can be adapted to downstream tasks with minimal additional supervision [1]. Among the most promising applications is industrial defect

segmentation, where the goal is to precisely delineate regions of a manufactured component that deviate from quality standards. Traditional approaches to this problem have relied on two contrasting paradigms: hand-crafted feature extractors paired with classical machine learning classifiers, and end-to-end supervised deep learning networks trained on large collections of labeled defect images [2][3]. Both methods suffer from critical limitations in dynamic production environments. Hand-crafted features often fail to capture subtle or novel defect morphologies, while supervised deep learning models require expensive annotation campaigns and degrade when the test distribution shifts away from the training distribution, a common occurrence when production lines introduce new materials, lighting conditions, or defect types [4].

Promptable vision foundation models address these limitations by offering a flexible interface through which a user can specify the segmentation target using a prompt, typically in the form of points, bounding boxes, or natural language descriptions [5]. The promptability mechanism allows the same model to segment a wide variety of defect classes without retraining, simply by providing appropriate prompts. This paradigm substantially reduces the need for large labeled datasets and enables rapid adaptation to new inspection tasks, making it particularly attractive for small- or medium-sized manufacturers that lack the resources to curate extensive defect databases [6]. However, the integration of such models into operational inspection systems introduces a host of system-level challenges that extend beyond the algorithmic core. Issues of latency, throughput, computational cost, data governance, model robustness, and fairness must be carefully addressed to ensure that the theoretical flexibility of promptable models translates into reliable, scalable, and equitable industrial deployments [7].

This paper adopts a systems perspective to analyze the deployment of promptable vision foundation models for industrial defect segmentation. We examine architectural choices that modulate the trade-off between general-purpose capability and specialized performance, infrastructure considerations for real-time or near-real-time inference, and governance frameworks needed to maintain quality assurance standards. We also explore the robustness of these models under domain shift and adversarial conditions, the fairness implications of their training data, and the sustainability costs associated with large-scale transformer architectures. Through cross-domain comparisons with established inspection methods, we identify the contexts in which promptable models offer decisive advantages and those where traditional approaches remain preferable. The paper concludes with a forward-looking discussion of research and policy directions, emphasizing the need for lightweight deployment strategies, continuous learning mechanisms, and regulatory alignment in safety-critical quality inspection.

2. Background and Evolution of Industrial Quality Inspection

Industrial quality inspection has evolved from manual visual checks conducted by human inspectors to increasingly automated systems based on machine vision. Early automated inspection relied on rule-based image processing techniques such as thresholding, edge detection, and morphological operations to segment defects [8]. These methods were computationally efficient but brittle, failing when variations in lighting, texture, or part geometry exceeded the assumptions embedded in hand-crafted rules. The introduction of convolutional neural networks (CNNs) in the 2010s marked a significant improvement, as end-to-end learning could discover hierarchical features directly from pixel data without manual engineering [9]. CNN-based defect segmentation, often implemented as U-Net or Mask R-CNN variants, achieved high accuracy on benchmark datasets and was widely

adopted in industries such as printed circuit board inspection, steel surface inspection, and textile quality control [2][3].

Despite these successes, CNN-based systems present substantial maintenance burdens. Each new defect type or product variant typically requires the collection of hundreds to thousands of annotated images, followed by retraining or fine-tuning of the entire model. This process is time-consuming and costly, and it creates a bottleneck in industries with high product mix and frequent design changes [10]. Moreover, the performance of supervised CNNs degrades under distributional drift, which can occur due to changes in ambient lighting, camera calibration, or subtle shifts in manufacturing processes. Retraining cycles must be carefully scheduled and validated, adding further operational complexity [4].

Vision foundation models, particularly those based on transformer architectures trained on web-scale datasets, have emerged as a powerful alternative. The Segment Anything Model (SAM) introduced a promptable segmentation paradigm that can segment any object in an image given a point, box, or mask prompt, with zero-shot generalization to unseen object categories [5]. Simultaneously, vision-language models such as CLIP and ALIGN enabled zero-shot classification and segmentation through natural language descriptions, bridging visual and semantic understanding [11][12]. These models represent a shift from task-specific training to task-agnostic pre-training, where the model acquires a rich representation of visual concepts that can be invoked via prompts. For industrial defect segmentation, this means that a single foundation model can potentially handle a wide variety of defect types, from cracks and scratches to corrosion and misalignment, without requiring defect-specific training data [6].

The term "promptable" captures the interactive nature of these models: the user provides a prompt that specifies the object of interest, and the model returns a segmentation mask. Promptability enables a human-in-the-loop workflow where inspectors can quickly annotate defects or correct model outputs by providing additional prompts, thereby creating a feedback mechanism that improves performance over time without full retraining [13]. This property is especially valuable in quality inspection, where the diversity of defects is large and continuously evolving. However, the benefits of promptability come with trade-offs in computational complexity, model size, and inference latency, which we analyze in the following sections.

3. Architecture of Promptable Vision Foundation Models

The architectural design of promptable vision foundation models determines their capacity for rapid few-shot adaptation and their suitability for industrial real-time requirements. Most promptable models follow an encoder-decoder structure where an image encoder processes the input image into a dense feature map, and an interactive decoder uses the prompt to generate a segmentation mask. SAM, for instance, employs a MAE-pre-trained ViT-Huge image encoder, a prompt encoder that handles points, boxes, and masks, and a lightweight mask decoder that outputs multiple candidate masks [5]. The prompt encoder injects information about the user-specified region, and the decoder attends to the corresponding features in the image embedding, enabling precise segmentation even for ambiguous boundaries.

Vision-language models such as CLIP adopt a dual-encoder architecture: one for images and one for text, with a shared embedding space learned via contrastive pre-training on image-text pairs [11]. For segmentation tasks, CLIP can be combined with pixel-level mapping networks

to produce segmentation masks given a language description of the defect. This approach is particularly powerful for industrial settings because it allows inspectors to specify defects in natural language, such as "scratch on metal surface" or "crack near weld joint", without needing to provide geometric prompts [14]. However, the spatial resolution of CLIP-based segmentation is often coarser than that of dedicated segmentation models, and fine-tuning may be required to achieve pixel-level accuracy [12].

A key architectural trade-off exists between model generality and specialization. Large transformer models with hundreds of millions or billions of parameters achieve high zero-shot performance across diverse domains, but they require substantial computational resources for inference. In contrast, smaller models with fewer parameters can be deployed on edge devices but may exhibit lower accuracy on novel defects or require more prompt adjustments to achieve acceptable performance [15]. The choice of architecture therefore depends on the specific constraints of the production environment: high-throughput assembly lines may prioritize fast inference on GPU-based edge nodes, while low-volume inspection of complex parts may tolerate cloud-based inference with higher latency.

Another architectural consideration is the ability to handle multi-modal prompts. In practice, industrial inspectors may combine point clicks with a verbal description, or provide a rough bounding box together with a semantic label. Some recent architectures, such as Grounding DINO, integrate detection and segmentation with language grounding, enabling prompts that are both spatial and semantic [16]. This flexibility can reduce the number of interactions required to achieve accurate segmentation, which is critical in time-sensitive inspection workflows. However, the added complexity of multi-modal processing increases the risk of inconsistencies when prompt modalities conflict, necessitating robust fusion mechanisms and careful user interface design [17].

4. System-Level Deployment and Infrastructure Considerations

Deploying promptable vision foundation models in industrial environments involves navigating a complex space of infrastructure choices, including hardware acceleration, network bandwidth, data storage, and software stack compatibility. The computational demands of transformer-based image encoders are substantial; for example, SAM's ViT-Huge encoder requires approximately 2.5 GFLOPs for a single 1024x1024 image, and the forward pass on a high-end GPU takes on the order of hundreds of milliseconds [5]. For production lines with cycle times under one second, this level of latency may be unacceptable unless the models are heavily optimized or the image resolution is reduced. Quantization, pruning, and knowledge distillation are common techniques to reduce model size and accelerate inference, but they can introduce accuracy degradation that must be quantified for each defect class [15].

The choice between on-premise and cloud deployment carries significant implications for data governance and latency. On-premise deployment, often on edge servers or dedicated GPU workstations, minimizes data transfer delays and ensures that sensitive manufacturing images remain within the factory network, addressing privacy and intellectual property concerns [7]. However, on-premise hardware must be procured, maintained, and upgraded, and it may not be able to support the largest foundation models without substantial capital expenditure. Cloud deployment offers elastic scalability and access to state-of-the-art hardware, but it introduces network latency, data egress costs, and dependency on internet connectivity. For many manufacturers, a hybrid approach is preferred: lightweight promptable models run on edge devices for real-time defect screening, while more powerful models in the cloud are invoked for ambiguous or novel defects that require human verification [18].

Infrastructure governance is another critical dimension. The deployment of AI-based inspection systems must comply with industry standards for quality management, such as ISO 9001 and ISO 13485 in medical device manufacturing. This requires rigorous validation of the model's performance under the expected operating conditions, including sensitivity and specificity for each defect type [19]. Promptable models, because they can be adapted on the fly via prompt changes, introduce a new challenge: the operational definition of a "defect" may shift with different prompters, potentially leading to inconsistent quality judgments. To mitigate this, organizations must establish strict protocols for prompt templates, define allowable prompt variations, and implement audit trails that record the prompts used for each inspected part [20]. Such governance measures ensure that the flexibility of promptable models does not undermine reproducibility and traceability, which are essential for quality assurance in regulated industries.

5. Robustness, Fairness, and Governance in Inspection Systems

Robustness is a paramount concern for any AI system deployed in safety-critical contexts. Promptable vision foundation models are trained on large, diverse datasets that may include images of natural scenes, everyday objects, and synthetic data, but industrial defects exhibit specific visual characteristics that are often underrepresented in those datasets. For example, micro-cracks in semiconductor wafers or delamination in composite materials have subtle texture and contrast properties that a model trained on generic internet images may not have learned to segment reliably [6]. Domain shift from the pre-training distribution to the industrial application domain can lead to systematic failures, where certain defect types are consistently missed or mis-segmented. Robustness can be improved through targeted fine-tuning on a small number of industrial images, but this partially negates the advantages of zero-shot promptability [5].

Adversarial robustness is also relevant in industrial settings, particularly in high-value manufacturing where sabotage or fraud could be attempted. Promptable models are vulnerable to adversarial perturbations of the input image that cause the prompt to be ignored or misinterpreted [21]. For instance, a small, imperceptible change in pixel values could make a crack appear as a harmless surface texture, leading the model to output a false negative. Defending against such attacks requires adversarial training or input preprocessing, but these measures add computational overhead and may reduce accuracy on benign examples. Moreover, the interactive nature of promptable models introduces a new attack surface: an adversary could manipulate the prompt itself (e.g., by feeding misleading coordinate points) to cause the model to segment a defect in a non-defective region. Governance frameworks must therefore include procedures for input validation and prompt verification [20].

Fairness considerations emerge from the composition of the pre-training datasets. If a foundation model is predominantly trained on images from certain geographic regions, lighting conditions, or product types, its performance may vary systematically across different manufacturing sites or material categories [22]. For instance, a model trained largely on images from automobile assembly in North America may perform poorly on textile defects in South Asian factories due to differences in texture, color, and lighting. This disparity can lead to inequitable quality outcomes, where certain production lines receive inferior inspection accuracy. Addressing fairness requires auditing the model's performance across diverse subgroups, collecting representative data from all deployment sites, and possibly using prompt engineering to calibrate outputs per site [23]. In regulated industries, fairness concerns

intersect with liability: if a model systematically misclassifies defects in a manner correlated with region or supplier, the manufacturer may face legal exposure.

Governance of promptable models also involves managing the drift of the underlying model. Foundation models are periodically updated by their developers, and such updates can change the segmentation behavior even for the same prompts. Without careful version control, an inspection system that was validated with one model version may produce different results after an upgrade, disrupting quality records and raising regulatory compliance issues. Therefore, organizations must freeze model versions, maintain thorough documentation of the model lineage, and re-validate performance after any model update [19]. Additionally, the interactive prompt mechanism means that human operators may develop idiosyncratic prompting styles, introducing unwanted variability. Standardizing prompts through templates and training operators reduces such variability, but it also limits the flexibility that makes promptable models attractive.

6. Sustainability and Scalability of Foundation Model-Based Quality Inspection

The environmental sustainability of large-scale AI systems is a growing concern, and promptable vision foundation models are among the most computationally intensive architectures in use. Training a model like SAM with a ViT-H encoder consumes thousands of GPU-hours and emits significant carbon dioxide equivalents [5]. While the training cost is amortized over many downstream applications, the inference cost per image remains high compared to smaller task-specific models. For a quality inspection system processing millions of parts per year, the aggregate energy consumption of running a large transformer encoder on each image can be substantial. This energy footprint conflicts with corporate sustainability goals and may be subject to future carbon taxes or regulatory limits [24].

Strategies for reducing the sustainability impact include model distillation, where a smaller student network is trained to mimic the behavior of the large foundation model on the specific domain of interest. Distilled models can retain high accuracy on the target defect classes while requiring an order of magnitude fewer computations [15]. Another approach is to use the foundation model as a gatekeeper: a lightweight classifier first determines whether a part likely contains a defect; only if the classifier flags the region does the full foundation model segment it accurately. Such cascaded architectures trade off worst-case latency for average efficiency. Additionally, hardware accelerators such as edge TPUs or FPGA-based inference engines can execute quantized models with much lower energy per inference, making on-premise deployment more sustainable [7].

Scalability of promptable model deployment involves not only computational resources but also the ability to manage an evolving catalog of prompts. As new product lines are introduced, operators must define and validate prompts for each defect type. Without a systematic prompt management system, the number of prompts grows quickly, and the risk of prompt conflict or inefficiency increases. Semantic prompt libraries, where prompts are grouped by defect family (e.g., "surface scratch", "edge crack", "color inconsistency") and validated against benchmark images, can improve consistency. However, developing such libraries requires cross-functional collaboration between domain experts and machine learning engineers, and it represents a non-trivial upfront investment [13].

From a scalability perspective, promptable models offer a significant advantage in handling low-frequency defects. In traditional supervised systems, rare defects are underrepresented in training data, leading to poor performance. With promptable models, a single prompt can be

designed for a rare defect type, and the model’s generalization ability may allow it to segment that defect even from few examples. This capability reduces the need for costly data collection campaigns for every rare defect class, thereby improving the economic scalability of quality inspection across a wide range of products [6].

7. Case Illustrations and Cross-Domain Comparisons

To contextualize the trade-offs discussed, we consider illustrative cases from three manufacturing sectors: electronics, automotive, and textile inspection. In electronics, printed circuit board (PCB) inspection requires detecting microscopic soldering defects, trace cracks, and component misplacement. Traditional CNN-based systems achieve high throughput with specialized networks that are optimized for the small field of view and high contrast of PCB images [2]. Deployment of SAM for PCB inspection revealed that while the model could segment gross defects such as missing components, it struggled with subtle soldering voids because the pre-training data did not include such fine-scale industrial imagery [5]. Fine-tuning on a small set of PCB images improved performance, but the inference latency of the full ViT-H model (approximately 0.5 seconds per image on an A100 GPU) was too high for a line running at 60 boards per minute. A distilled smaller model, though faster, reduced accuracy on the most challenging defect classes [15]. This case illustrates the tension between generality and speed: promptable models excel in flexibility but require careful optimization for high-speed production.

In automotive body panel inspection, defects include dents, scratches, and paint imperfections that vary in size and contrast. Manufacturers have adopted promptable models as a tool for human-in-the-loop inspection, where an operator quickly clicks on a potential defect and the model segments it, then the operator verifies or corrects the mask. This workflow reduces the cognitive load on inspectors and speeds up the decision process [13]. Unlike fully automated systems, the loop tolerates higher latency (a few seconds per part) because the operator’s actions are the bottleneck. Here, the promptable model’s ability to handle ambiguous boundaries and partial occlusions (common in curved panel surfaces) makes it superior to traditional edge-based algorithms. The trade-off is that operator prompting consistency becomes a source of variability, requiring periodic calibration.

Textile inspection presents a different challenge: defects such as weft bars, stains, and holes often appear against highly textured backgrounds. Foundation models trained on natural images tend to treat repetitive textile patterns as features of interest, leading to false positives. However, using language prompts like "stain on fabric" combined with a point prompt reduces false detections because the language priors guide the model away from natural texture variations [14]. This cross-modal prompting demonstrates the advantage of vision-language models over vision-only models. The case also highlights the importance of domain-specific prompt design, which requires collaboration with textile engineers to create unambiguous defect descriptions.

Cross-domain comparisons with traditional machine vision show that promptable models are most beneficial when defect types are numerous, infrequent, or evolving. For high-volume production with a stable defect portfolio, a dedicated task-specific CNN remains more efficient in terms of speed, cost, and energy. Conversely, for small-batch production, custom manufacturing, or research and development lines, the flexibility of promptable models outweighs their higher per-inference cost. The optimal deployment strategy often involves a hybrid system, where a lightweight CNN acts as a pre-filter and a promptable foundation model handles only the ambiguous cases.

8. Future Directions and Policy Implications

The trajectory of promptable vision foundation models for industrial defect segmentation points toward greater specialization through domain-specific fine-tuning, more efficient architectures, and tighter integration with digital twin and monitoring systems. Lightweight variants of SAM, such as MobileSAM and EfficientSAM, have already demonstrated that strong performance can be achieved with an order of magnitude fewer parameters, making real-time inference on edge devices feasible [15][5]. Future research should focus on continuous adaptation mechanisms that can update the model's behavior in response to gradual concept drift without full retraining, perhaps through online prompt optimization or memory-augmented decoders.

Policy implications span multiple dimensions. First, the certification of AI-based inspection systems by regulatory bodies such as the Food and Drug Administration for medical devices or the Federal Aviation Administration for aerospace components will require clear guidelines on promptable models. How does one certify a system whose behavior depends on user prompts? One possibility is to certify the model-prompt combination as an integrated system, where prompts are part of the validated configuration. Second, data privacy regulations, including the General Data Protection Regulation and the California Consumer Privacy Act, may impose restrictions on sending manufacturing images to cloud-based model inference services, especially if the images contain proprietary design information. On-premise deployment or federated learning approaches could mitigate these concerns while still benefiting from foundation model capabilities [24]. Third, the environmental footprint of large-scale AI must be accounted for in corporate sustainability reporting, incentivizing the development of energy-efficient model architectures and hardware.

Finally, the fairness and equity of inspection quality across global supply chains remains an underexplored area. Manufacturers in developing countries may lack the infrastructure to deploy large foundation models, potentially widening the quality gap between high-tech and low-tech production sites. Open-source promptable models, low-cost edge hardware, and capacity-building initiatives could help democratize access, but these efforts require coordinated action from industry consortia, governments, and academic institutions [22].

9. Conclusion

Promptable vision foundation models represent a transformative approach to industrial defect segmentation and quality inspection, offering unprecedented flexibility through few-shot and zero-shot adaptation. This paper has provided a system-level analysis of their deployment, encompassing architectural trade-offs, infrastructure requirements, robustness and fairness considerations, and sustainability implications. The ability to segment a wide array of defect types using simple point, box, or language prompts reduces the dependency on large annotated datasets and enables rapid response to changing production conditions. However, the computational cost of large transformer models, the need for careful governance of prompts and model versions, and the challenges of domain shift and adversarial robustness demand a holistic perspective that goes beyond algorithmic performance. Cross-domain case studies illustrate that promptable models excel in contexts with high defect diversity and low volume, while traditional methods remain advantageous for high-speed, stable production lines. Future research should prioritize efficient model variants, continuous adaptation strategies, and policy frameworks that ensure safe, fair, and sustainable deployment across global manufacturing ecosystems.

References

1. Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., von Arx, S., ... & Liang, P. (2021). On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*.
2. Ren, R., Hung, T., & Tan, K. C. (2017). A comprehensive survey of surface defect detection. *IEEE Transactions on Industrial Informatics*, 14(2), 415-428.
3. Kong, L., & Niu, Y. (2022). Deep learning-based defect detection in industrial manufacturing: A survey. *Journal of Manufacturing Systems*, 62, 541-559.
4. Zhang, J., & Li, X. (2020). Domain adaptation for industrial defect detection: A critical review. *IEEE Transactions on Industrial Electronics*, 68(9), 8567-8578.
5. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., ... & Girshick, R. (2023). Segment anything. In *Proceedings of the IEEE/CVF International Conference on Computer Vision* (pp. 4015-4026).
6. Li, Y., Mao, H., & Girshick, R. (2023). Exploring the potential of segment anything model for industrial defect segmentation. *arXiv preprint arXiv:2304.07849*.
7. Wu, C., Li, X., & Zhou, B. (2022). Edge-cloud collaboration for industrial AI: Architecture and challenges. *IEEE Network*, 36(6), 56-63.
8. Malamas, E. N., Petrakis, E. G. M., Zervakis, M., Petit, L., & Legat, J. D. (2003). A survey on industrial vision systems, applications and tools. *Image and Vision Computing*, 21(2), 171-188.
9. Simonyan, K., & Zisserman, A. (2014). Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*.
10. Wang, Z., & Liu, Y. (2021). Active learning for defect detection in small-batch manufacturing. *IEEE Transactions on Automation Science and Engineering*, 19(3), 1456-1467.
11. Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., ... & Sutskever, I. (2021). Learning transferable visual models from natural language supervision. In *Proceedings of the 38th International Conference on Machine Learning* (pp. 8748-8763).
12. Li, J., Li, D., Xiong, C., & Hoi, S. (2022). BLIP: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *Proceedings of the 39th International Conference on Machine Learning* (pp. 12888-12900).
13. Cheng, T., Wang, X., Chen, L., & Liu, W. (2023). Interactive image segmentation with SAM: A human-in-the-loop evaluation for industrial inspection. *arXiv preprint arXiv:2305.12345*.
14. Yan, Z., & Zhang, H. (2023). Zero-shot industrial defect segmentation using vision-language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops* (pp. 4123-4132).
15. Zhao, W., Hu, J., & Chao, H. (2023). Efficient segment anything: Distilling SAM for real-time industrial applications. *arXiv preprint arXiv:2305.08727*.

16. Liu, S., Zeng, Z., Ren, T., Li, F., Zhang, H., Yang, J., ... & Lin, D. (2023). Grounding DINO: Marrying DINO with grounded pre-training for open-set object detection. arXiv preprint arXiv:2303.05499.
17. Diao, H., & He, T. (2023). Robust multi-modal prompting for industrial defect segmentation. In Proceedings of the 31st ACM International Conference on Multimedia (pp. 1201-1210).
18. Jiang, J., & Xu, J. (2022). Hybrid edge-cloud architecture for real-time industrial defect detection. IEEE Transactions on Cloud Computing, 10(4), 2145-2158.
19. International Organization for Standardization. (2015). Quality management systems – Requirements (ISO 9001:2015).
20. Raji, I. D., Smart, A., White, R. N., Mitchell, M., Gebru, T., Hutchinson, B., ... & Barnes, P. (2020). Closing the AI accountability gap: Defining an end-to-end framework for internal algorithmic auditing. In Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency (pp. 33-44).
21. Madry, A., Makelov, A., Schmidt, L., Tsipras, D., & Vladu, A. (2018). Towards deep learning models resistant to adversarial attacks. In International Conference on Learning Representations.
22. Buolamwini, J., & Gebru, T. (2018). Gender shades: Intersectional accuracy disparities in commercial gender classification. In Proceedings of the 1st Conference on Fairness, Accountability and Transparency (pp. 77-91).
23. Chouldechova, A., & Roth, A. (2020). A snapshot of the frontiers of fairness in machine learning. Communications of the ACM, 63(5), 82-89.
24. Strubell, E., Ganesh, A., & McCallum, A. (2019). Energy and policy considerations for deep learning in NLP. In Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (pp. 3645-3650).