Computer Vision-Based In-Line Inspection in Intelligent Quality Control Systems: Methods, Applications, and Engineering Challenges

Peng Minghui
Hebei University of Architecture and Engineering, China

ABSTRACT

On-line visual inspection-that is, automatically taking images and analyzing them during manufacturing, rather than after manufacturing-is very important for intelligent quality control of modern factories. It replaces slow and subjective manual inspection with rapid and objective detection. This paper reviews the on-line industrial inspection method based on computer vision and deep learning, using the theory of network physical system (CPS), system engineering and information theory. Inspection chain is regarded as a channel to obtain information, and its capacity must be large enough to track production and find defects. It lays a theoretical foundation for intelligent quality control system with vision, which is regarded as closed-loop CPS, in which the fault signal from the image always adjusts the machine settings, not just that the product is defective. We analyzed four changes: from sample inspection to continuous whole surface inspection, from once discarding to predicting defects before they arrive, from manual judgment to closed-loop automatic quality control, and from single inspection station to integrated quality -CPS ecosystem. The applications in three main fields-surface defect detection of metal materials, unattended anomaly detection when defects are rare, and real-time quality control of edge equipment-are illustrated by the reference data set (MVTec AD), the latest architecture and factory use examples. There are still thorny problems such as class imbalance, the cost of tagging data, the difficulty of generalization from one domain to another, and the choice between speed and accuracy on edge devices. A new framework concept puts forward that the uncertainty of detection is regarded as an important signal to improve the production process, and the quality control is changed from simple "good or bad" to a means to continuously improve manufacturing.

Keywords: Computer Vision; In-Line Inspection; Defect Detection; Anomaly Detection; Quality Control; Cyber-Physical Systems; Edge Computing; Deep Learning.

1. INTRODUCTION

Quality control in manufacturing serves the dual function of filtering non-conforming products from the output stream and generating diagnostic information about the processes producing them. Traditional human visual inspection performs both functions imperfectly: it is susceptible to fatigue-induced missed detections, inherently limited in throughput, incapable of generating spatially precise, quantitative defect records, and structurally disconnected from the process-control layer that could act on its findings [1]. These restrictions become more and more important when the products are more complex, with fewer defects, and the manufacturing speed exceeds the speed that humans can check. This is the situation of manufacturing semiconductors, steel, textiles and precision mechanical parts now.

Computer vision has gradually replaced human beings to inspect products. It starts with hand-made textures and gradients, then a simple classifier, and now a deep convolutional neural network. These networks can learn by themselves how to find defects in images without telling them how to do it [2]. In this evolution, three things are very important to the industry. First of all, we created an object detection architecture, which can find and classify multiple defects at the same time, and the speed is very fast, almost as fast as the production speed [2]. Then unsupervised and self-supervised anomaly detection methods are invented, which only need defect-free images to learn, which is convenient because there are few enough defect samples in the factory [3,4]. Finally, the edge computing platform has become powerful enough to run the neural network near the sensor without waiting for the central server, which will slow down the rapid production chain.

This review studies these developments under the framework of CPS theory and system engineering. She said that their comprehensive effect is to transform the quality visual inspection of an isolated measurement function into a feedback node, which is located in a larger CPS that controls production quality. The rest is like this. Section 2 develops the theoretical basis. The third part analyzes four transformation paths driven by AI. Section 4 focuses on the application of surface defect detection, unsupervised anomaly detection and quality control of edge deployment. Section 5 critically examines the challenges and prospects. End of section 6

2. THE THEORETICAL FOUNDATION OF ARTIFICIAL INTELLIGENCE AND INTELLIGENT ENGINEERING SYSTEMS

2.1 The Connotation and Characteristics of Intelligent Engineering Systems

In the quality inspection paradigm, intelligent engineering system is a network physical system, in which the visual state of manufacturing objects (images taken at production speed) is always analyzed to update the probability model of product quality, and then the model is returned to the process control layer [5]. There are three things that make these systems different from ordinary automatic inspection machines (called AOI). First of all, they have been working in real time. They don't just look at a few parts, but look at all the parts at the speed of the machine. They make a comprehensive quality report on everything they make, not just an estimate. Secondly, they know exactly where the defect is and what it looks like. Modern systems using deep learning can not only tell us that there are defects, but also tell us where it is and what type it is. This information helps to understand why the failure occurred. Third, they can send messages to change the settings of the machine. When the quality signal is connected to the machine control, the defect information can adjust the parameters in real time, instead of just discarding the wrong parts, which changes the inspection: it is no longer the filter at the end, but the sensor in the manufacturing process [5].

2.2 The Theoretical Logic of Artificial Intelligence Empowering Engineering Systems

From an information-theoretic perspective, the in-line visual inspection pipeline can be modeled as a multi-stage channel: the acquisition sub-channel (camera, lighting, optics) converts physical surface state into raw pixel data; the feature-extraction sub-channel (convolutional layers) compresses raw pixel data into a defect-relevant representation; and the classification or anomaly-scoring sub-channel maps that representation to a detection decision. The capability of this channel-mutual information between surface state and output detection-is limited by the spatial resolution and signal noise characteristics of subchannel acquisition, and is discriminated by subchannel classification at the top [2]. Deep learning improves the ability of subchannel classification by learning task adaptation features directly from data, instead of accepting the upper limit imposed by manual engineering descriptors. However, transmission learning from large-scale pre-trained convolutional networks further extends this ability to defect identification tasks, in which the lack of labeled training data will prevent convergence from Scratch [6].

From the point of view of system engineering, the important question is not only whether the pipeline inspection has achieved good sensitivity and specificity on the benchmark data set, but also whether its output is formatted and has enough spatial accuracy, semantic richness and time sensitivity to help reject decision-making and process correction inference. A detection architecture that outputs a binary pass/fail signal is complete for the rejection function but discards the localization and classification information necessary for the process-correction function; the theoretical case for in-line vision in intelligent engineering systems depends on deploying architectures that preserve this richer signal through to the quality-control and process-optimization layers.

3. THE TRANSFORMATION PATH OF INTELLIGENT ENGINEERING SYSTEMS MODE DRIVEN BY ARTIFICIAL INTELLIGENCE

3.1 Transition from Standardized Production to Personalized and Flexible Manufacturing

Standardized mass production permitted inspection systems calibrated once for a single product configuration; flexible manufacturing requires inspection adaptable to rapidly changing product variants without manual recalibration. Deep-learning inspection architectures address this requirement through transfer learning: a backbone pre-trained on large general image datasets provides a rich, general-purpose feature hierarchy that can be fine-tuned for a new product variant using only a small number of labeled examples, reducing the recalibration cost from days of data collection to hours of fine-tuning [6]. Unsupervised anomaly detection methods extend this adaptability further, since models trained exclusively on defect-free images of the new variant require no defect labels at all for adaptation—only normal-state reference images, which are readily available for any product configuration that has been manufactured [3, 4].

3.2 Transition from Reactive Response to Proactive Prediction and Optimization

Reactive quality control identifies defects after they are produced and rejects the affected products; proactive quality control anticipates the process conditions under which defects will emerge and intervenes before they materialize. In-line vision enables this transition by providing a high-density temporal record of defect occurrence that can be analyzed for systematic patterns predictive of impending process drift. When a convolutional detection model reports spatially clustered surface scratches in steel strip output, for example, that spatial pattern is diagnostic of roll wear rather than random process noise; a quality-CPS that routes this pattern back to the rolling-mill control layer can trigger roll exchange before the defect rate exceeds specification, rather than after product has been scrapped [7]. This transition from reactive to proactive is contingent on the inspection system's ability to produce semantically typed, spatially localized defect records rather than undifferentiated rejection counts, reinforcing the theoretical argument of Section 2.2 for architectures that maximize output semantic richness.

3.3 Transition from Manual Operation to Autonomous Intelligent Control

Full autonomous quality control closes the loop completely: inspection results parameterize process-control actuators without human intermediation. Achieving this requires not only high detection accuracy but also quantified uncertainty: an autonomous process-adjustment system must know how confident the inspection model is in each classification, so that borderline detections trigger conservative interventions rather than aggressive actuation that could destabilize the process. Uncertainty quantification in deep-learning inspection models—through Bayesian neural network approximations, ensemble disagreement, or conformal prediction—is therefore not a research nicety but a functional prerequisite for genuinely autonomous quality control. Edge-computing architectures are equally essential: autonomous process adjustment requires that the defect signal reach the actuator within the process response time, which at high-speed production rates is measured in milliseconds and is incompatible with round-trip latency to centralized inference servers.

3.4 Transition from Isolated Systems to Integrated Cyber-Physical Ecosystems

Isolated detection stations produce high-quality data. These data are often recorded for later search, but they are almost never used in real time by upstream process controllers, maintenance systems or scheduling engines. When you integrate all these into the network physical quality ecosystem, the situation will change: the defect types that exceed a certain statistical threshold will trigger the maintenance suggestions of related tools, send information to the parameter optimizer process, and tell the scheduling system to put the affected equipment in a pending state. In order to make this integration work, we need a shared semantic representation of defect types and causality in the process-an ontology layer, which can transform the morphological classification of defects by visual system into process engineering that can be understood and used by upstream executors and scheduling agents [5]. We are improving the existing transformation framework, suggesting that testing trust should be allocated to testing systems instead of simple decision-making, which is the main signal spread in the quality ecosystem. This allows the upstream system to measure its response according to the certainty of the fault, thus avoiding the traditional typical binary decision-making based on pass/fail check and critical value decision-making.

4. APPLICATIONS OF ARTIFICIAL INTELLIGENCE IN INTELLIGENT ENGINEERING SYSTEMS

4.1 Application in Intelligent Fault Diagnosis and Predictive Maintenance

Visual surface inspection and fault diagnosis converge in applications where surface morphology is a direct indicator of tool or machine condition. For flat steel production—hot-rolled and cold-rolled strip being among the highest-throughput vision-inspection environments—comprehensive reviews of automated visual defect detection demonstrate that convolutional architectures trained on categories including scale pits, scabs, cracks, and roll prints consistently outperform texture-descriptor-based classical methods in both detection rate and false-alarm rate when sufficient training data is available, while simultaneously generating the spatial and semantic defect records needed to support downstream roll-wear diagnostics [7]. The surface defect morphology carries direct diagnostic information: elongated, periodically spaced marks at a spatial period matching the roll circumference indicate roll wear rather than random metallurgical flaw, and routing this diagnosis to the maintenance system allows predictive roll replacement rather than reactive response to out-of-tolerance product. This diagnostic coupling is only exploitable, however, if the inspection system outputs typed, localized detections rather than binary pass/fail counts, again emphasizing the importance of architecture choices that preserve semantic richness throughout the detection pipeline.

4.2 Application in Intelligent Process Optimization and Quality Control

Unsupervised anomaly detection methods have become the dominant approach to quality control in environments where defect samples are too few or too diverse to support supervised training, which characterizes the majority of real industrial scenarios outside of high-volume commodities. The MVTec Anomaly Detection dataset—comprising 5,354 high-resolution images of fifteen object and texture categories with pixel-precise anomaly ground truth annotations—established the benchmark against which unsupervised methods based on reconstruction (convolutional autoencoders, variational autoencoders, generative adversarial networks) and feature-embedding (pre-trained backbone descriptors, memory banks, normalizing flows) are evaluated, and initial results revealed substantial performance gaps between these methods and human inspection performance, motivating the extensive subsequent literature on unsupervised industrial inspection [4]. The recent systematic investigation on unsupervised anomaly detection of industrial images shows that the feature embedding method using pre-trained convolution backbone descriptors and neighbor anomaly scores is always superior to the reconstruction-based method in various types of industrial defects. In addition, when they are used, their calculation cost is lower, which makes them easier to put on the machines on the production line [3]. As a result of these advances, it is now possible to use quality control systems to check defects that could not be checked before, just by giving normal training images. This eliminates the problem of annotation. Previously, only products with a large number of products and a large number of defects were subject to supervision and economic inspection.

4.3 Application in Intelligent Human-Machine Collaboration and Operational Efficiency Enhancement

There are two forms of man-machine cooperation in vision-based quality control. The first is active collaboration during model development, in which the annotation interface-weakly supervised and semi-supervised training paradigm-which allows quality engineers to interactively confirm or correct model detection greatly reduces labeling work compared with pixel-by-pixel comprehensive annotation, while maintaining the detection performance close to the fully supervised benchmark [8]. The second is cooperation in the production process. The model shows its detection and uncertainty-salinity map, thermal map of pixel-level anomaly and defect examples from reference library to quality engineers through easy-to-understand images. In this way, when the model is not clear, human can decide to change the model decision without the operator personally checking each product. The research on the unsupervised deep learning method for finding anomalies shows that finding anomalies of each pixel is the most helpful for this second form of cooperation, because the operator can quickly verify the detection by only looking at the marked area instead of the whole product [8]. By observing how people use these collaborations in real life, we can see that when the model works with human beings-the model decides by itself in a very safe situation and leaves the difficult situation to human beings-we get more finished products and fewer defects, instead of doing all the work by hand or using machines, especially when we start to manufacture new products, the model has not learned everything.

5. CHALLENGES AND PROSPECTS OF ARTIFICIAL INTELLIGENCE EMPOWERING INTELLIGENT ENGINEERING SYSTEMS

5.1 Current Main Challenges

Four challenges dominate current in-line vision-inspection practice. First, class imbalance between normal and defect samples remains endemic: defect events are by definition rare in a well-controlled process, creating severe training imbalance for supervised methods and making the false-alarm rate—rather than detection sensitivity—the binding performance constraint in high-throughput deployment, since even a modest false-alarm rate generates economically unacceptable volumes of falsely rejected product [7]. Second, cross-domain generalization is fragile: models trained on one product material, lighting condition, or defect morphology frequently fail to transfer to superficially similar settings without retraining, a fundamental limitation for flexible manufacturing contexts where product variety is high and the economic case for per-variant supervised training is weak [6]. Third, annotation cost for pixel-precise defect segmentation—the output format most useful for closed-loop process control—is orders of magnitude higher than image-level labeling, creating a systematic tension between the richness of annotation required by process engineers and the annotation budget that can realistically be sustained [3]. Fourth, edge-deployment constraints impose latency-accuracy trade-offs that are not yet systematically resolved: model compression, quantization, and pruning techniques that reduce inference latency on edge hardware also reduce detection sensitivity in ways that are highly dependent on defect scale, contrast, and frequency, and no standardized methodology for navigating this trade-off under specified production requirements currently exists [1].

5.2 Prospects for Future Development Trends

Three trajectories show the greatest promise. Foundation models and self-supervised visual representation learning, pre-trained on billions of diverse images, offer the prospect of zero-shot or few-shot defect detection without product-specific supervised training, potentially eliminating both the annotation cost and the cross-domain generalization barriers simultaneously, though their inference latency and computational requirements remain prohibitive for direct edge deployment without substantial model distillation [2]. Synthetic defect augmentation through physics-based rendering or generative adversarial networks—producing photorealistic defect images with controllable morphology inserted into defect-free backgrounds—offers a scalable path to addressing class imbalance and annotation scarcity, though the realism and domain coverage of synthetic defects relative to production defects remain active research questions [3]. Building on the inspection-uncertainty-as-process-signal extension proposed in Section 3.4, the most consequential long-term prospect may be the architectural integration of vision-inspection uncertainty into closed-loop process-control frameworks—not as a post-hoc reporting metric but as a real-time continuous input that parameterizes process adjustment magnitude and direction proportionally to detection confidence, transforming the inspection function from a discrete quality gate into a continuous, graded process-improvement actuator embedded within the production CPS.

6. CONCLUSION

This review has examined deep-learning-based computer vision as the mechanism through which artificial intelligence operationalizes in-line quality inspection within intelligent, cyber-physical manufacturing systems, grounding the analysis in CPS theory, systems engineering, and information theory. The central theoretical contribution is the framing of inspection architecture choices as information-channel design decisions: the semantic richness and spatial precision of detection outputs determine not only whether products can be correctly classified but whether the inspection function can contribute to process optimization and closed-loop quality control beyond the purely reactive rejection function it historically performed. Across the four transformation paths examined, increasing quality-system capability consistently demands richer, more semantically typed, more spatially precise detection outputs, met by complementary methodological advances including region-based deep detection, unsupervised feature-embedding anomaly scoring, and uncertainty-aware inference. Applications in surface defect detection, unsupervised anomaly detection, and human-machine collaborative inspection demonstrate the breadth of industrial environments in which these methods are achieving or approaching human inspection performance, while persistent challenges in class imbalance, cross-domain generalization, annotation cost, and edge latency constrain deployment scope. The proposed reframing of detection uncertainty as a continuous process-control input offers a concrete architectural path toward quality systems that improve process performance continuously rather than flagging product failures after the fact.

REFERENCES

[1] Czimmermann, T., Ciuti, G., Milazzo, M., Chiurazzi, M., Roccella, S., Oddo, C. M., & Dario, P. (2020). Visual-based defect detection and classification approaches for industrial applications-A survey. Sensors, 20(5), 1459. https://doi.org/10.3390/s20051459

[2] Zhao, Z.-Q., Zheng, P., Xu, S.-T., & Wu, X. (2019). Object detection with deep learning: A review. IEEE Transactions on Neural Networks and Learning Systems, 30(11), 3212-3232. https://doi.org/10.1109/TNNLS.2018.2876865

[3] Cui, Y., Liu, Z., & Lian, S. (2023). A survey on unsupervised anomaly detection algorithms for industrial images. IEEE Access, 11, 55297-55315. https://doi.org/10.1109/ACCESS.2023.3282993

[4] Bergmann, P., Fauser, M., Sattlegger, D., & Steger, C. (2019). MVTec AD-A comprehensive real-world dataset for unsupervised anomaly detection. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (pp. 9592-9600). https://doi.org/10.1109/CVPR.2019.00982

[5] Monostori, L., Kádár, B., Bauernhansl, T., Kondoh, S., Kumara, S., Reinhart, G., Sauer, O., Schuh, G., Sihn, W., & Ueda, K. (2016). Cyber-physical systems in manufacturing. CIRP Annals, 65(2), 621-641. https://doi.org/10.1016/j.cirp.2016.06.005

[6] Luo, Q., Fang, X., Liu, L., Yang, C., & Sun, Y. (2020). Automated visual defect detection for flat steel surface: A survey. IEEE Transactions on Instrumentation and Measurement, 69(3), 626-644. https://doi.org/10.1109/TIM.2019.2963555

[7] Bergmann, P., Batzner, K., Fauser, M., Sattlegger, D., & Steger, C. (2021). The MVTec anomaly detection dataset: A comprehensive real-world dataset for unsupervised anomaly detection. International Journal of Computer Vision, 129(4), 1038-1059. https://doi.org/10.1007/s11263-020-01400-4

[8] Tao, X., Gong, X., Zhang, X., Yan, S., & Adak, C. (2022). Deep learning for unsupervised anomaly localization in industrial images: A survey. IEEE Transactions on Instrumentation and Measurement, 71, 5018021. https://doi.org/10.1109/TIM.2022.3196436

[9] Tao, F., Zhang, H., Liu, A., & Nee, A. Y. C. (2019). Digital twin in industry: State-of-the-art. IEEE Transactions on Industrial Informatics, 15(4), 2405-2415. https://doi.org/10.1109/TII.2018.2873186

[10] Lu, Y., Liu, C., Wang, K. I. K., Huang, H., & Xu, X. (2020). Digital Twin-driven smart manufacturing: Connotation, reference model, applications and research issues. Robotics and Computer-Integrated Manufacturing, 61, 101837. https://doi.org/10.1016/j.rcim.2019.101837