

Review
**Deep Learning for Ceramic Surface
Defect Inspection: A State-of-the-Art Survey**

B.S.Wang¹, L. Zhang^{*2}

¹School of Computer and Software, Nanyang Institute of Technology, Nanyang, 473004, China.

²Changchun Oubang Biotechnology Co. Ltd., Changchun, 130000, China.

received April 14, 2025; received in revised form May 11, 2025; accepted May 20, 2025

Abstract

Surface defect detection in ceramic manufacturing is crucial for ensuring product quality and functional performance. This comprehensive review examines the evolution and current state of deep learning applications in ceramic surface inspection. Traditional inspection methods, including manual visual inspection and conventional image processing techniques, face limitations in handling complex textures, varying lighting conditions, and subtle defect patterns. The review analyzes various deep learning architectures, from basic Convolutional Neural Networks (CNNs) to sophisticated region-based networks and segmentation models, highlighting their strengths and limitations in ceramic defect detection. Implementation considerations, including data acquisition, preprocessing strategies, model training approaches, and real-time processing requirements, are discussed extensively. The paper addresses critical challenges such as limited labeled data availability, class imbalance, and the need for robust performance under varying production conditions. Recent innovations in network architectures, training methodologies, and deployment strategies that address industry-specific challenges are examined. The review provides valuable insights for both researchers and practitioners, bridging the gap between academic research and industrial implementation while identifying emerging trends and future research directions in this rapidly evolving field.

Keywords: Automated inspection, surface quality control, machine vision, neural networks, manufacturing automation

I. Introduction

The ceramic industry represents one of the oldest yet most technologically evolving manufacturing sectors, with applications ranging from traditional pottery to advanced technical ceramics used in electronics, healthcare, and aerospace. In this high-precision industry, surface quality plays a paramount role in determining both the aesthetic appeal and functional performance of ceramic products^{1–3}. Even minor surface defects can significantly compromise product integrity, leading to functional failures and substantial economic losses. This criticality has driven continuous innovation in inspection methodologies, with the latest frontier being the application of deep learning technologies for automated defect detection⁴. Surface defects in ceramic products manifest in various forms, including cracks, pinholes, spots, scratches, glaze deviations, and dimensional irregularities^{5,6}. The detection of these defects presents unique challenges due to the diverse nature of ceramic surfaces, varying lighting conditions, and the subtle nature of many defects that may be barely perceptible to the human eye. Traditional manual inspection, while still prevalent in many facilities, suffers from inherent limitations of human fatigue, subjective assessment, and inconsistent evaluation criteria^{7,8}. These

limitations, coupled with increasing production speeds and quality standards, have necessitated the development of automated inspection systems.

The importance of surface defect detection in the ceramics industry cannot be overstated. In technical ceramics, where components are often used in critical applications like medical implants or electronic substrates, surface integrity directly impacts functionality and reliability^{9,10}. For architectural ceramics such as tiles and sanitaryware, surface quality affects both aesthetic appeal and market value. Early detection of defects not only enables immediate corrective actions but also provides valuable feedback for process optimization¹¹. This feedback loop is crucial for implementing zero-defect manufacturing strategies and maintaining competitive advantage in an increasingly demanding market. The evolution of ceramic surface inspection methods mirrors the broader technological advancement in manufacturing quality control¹². Early automated systems relied on conventional image processing techniques, employing basic algorithms for edge detection, thresholding, and morphological operations¹³. While these methods proved effective for simple defect patterns, they often struggled with complex textures, varying lighting conditions, and subtle defect characteristics¹⁴. The introduction of traditional machine

* Corresponding author: koffnys2015@163.com

learning approaches brought improvements through feature engineering and statistical pattern recognition, but these methods still required extensive domain expertise for feature design and parameter tuning.

The paradigm shift toward deep learning represents a fundamental transformation in surface defect inspection. Deep learning approaches, particularly convolutional neural networks (CNNs) have demonstrated remarkable capabilities in automatically learning hierarchical features from raw image data, eliminating the need for manual feature engineering^{15,16}. This ability to learn complex patterns and generalize across varying conditions has made deep learning particularly suitable for ceramic surface inspection, where defect patterns can be intricate and environmental conditions variable¹⁷. Recent years have witnessed an explosion of research in applying deep learning to ceramic surface inspection, with architectures ranging from basic CNNs to sophisticated region-based networks and segmentation models¹⁸. These approaches have not only improved detection accuracy but also enabled real-time processing capabilities, making them suitable for inline inspection systems¹⁹. The success of deep learning in this domain has been facilitated by advances in computing hardware, particularly graphics processing units (GPUs), and the availability of larger datasets for training²⁰. However, the implementation of deep learning for ceramic surface inspection is not without challenges. The industry faces issues related to limited availability of labeled training data, class imbalance between defective and non-defective samples, and the need for robust performance under varying production conditions²¹. Additionally, the integration of deep learning systems into existing manufacturing processes requires careful consideration of real-time processing requirements, hardware infrastructure, and operator training needs²².

This comprehensive review aims to provide a systematic analysis of the current state of deep learning applications in inspection of ceramic surface defects. This review begins by examining traditional inspection methods and their limitations, establishing the context for the adoption of deep learning approaches. The subsequent sections delve into various deep learning architectures that have been successfully applied to ceramic inspection, analyzing their strengths, limitations, and specific application scenarios. This review also explores critical implementation considerations, including data preparation, training strategies, and deployment challenges. The review particularly focuses on recent innovations in network architectures, training methodologies, and practical deployment strategies that address industry-specific challenges. By examining both theoretical foundations and practical applications, this review aims to bridge the gap between academic research and industrial implementation. The final sections discuss emerging trends, ongoing challenges, and future research directions, providing a roadmap for both researchers and practitioners in this rapidly evolving field.

II. Traditional Surface Defect Detection Methods

Visual inspection has historically been the primary method for detecting surface defects in ceramic products.

This approach relies on trained human inspectors who visually examine each product for potential flaws, drawing upon their experience to identify and classify various types of defects. However, visual inspection faces several significant limitations that impact its effectiveness in modern manufacturing environments. Human inspectors are susceptible to fatigue, especially during long inspection sessions, which can lead to inconsistent detection rates throughout a work shift. Additionally, the subjective nature of human assessment often results in variability between different inspectors, making it difficult to maintain standardized quality criteria. The speed limitations of manual inspection also create bottlenecks in high-throughput production lines, where rapid assessment is crucial for maintaining production efficiency.

(1) Image processing techniques

The limitations of manual inspection led to the development of automated inspection systems based on digital image processing techniques. Although manual visual inspection is preferable in certain cases, it is inefficient and costly, requires a long training time, and causes fatigue. Moreover, in some industrial applications, manual visual inspection is not feasible because dangerous events may occur owing to errors in the inspection²³. Edge detection has been a fundamental approach in this domain, employing various algorithms to identify discontinuities in pixel intensity that might indicate surface defects. However, conventional feature extractors like a histogram of oriented gradients or local binary patterns (LBPs) combined with support vector machine (SVM) have shown limitations. For example, Song and Yan²⁴ improved the LBP method by combining it with an SVM, while Ghorai *et al.*²⁵ proposed an SVM approach based on wavelet features for defect classification tasks. However, these feature extractors can only analyze features in a single image, limiting their effectiveness for complex defect detection scenarios. More sophisticated approaches combining multiple imaging techniques and adaptive processing methods have emerged to address these limitations. For instance, Li *et al.*²⁶ proposed a defect detection method for strip steel production using the improved You Only Look Once (YOLO) network. While YOLO can classify defect images into six types and obtain position and size information, it faces challenges when dealing with subtle defects and complex backgrounds containing grain boundaries. As shown in Fig. 1, a more effective approach utilizes a combination of different lighting conditions (ring light, pseudo-coaxial light, and flat-dome light) to capture different aspects of surface characteristics. This multi-illumination setup facilitates better discrimination between actual defects and normal surface features by enhancing contrast, revealing surface topography, and minimizing shadow-induced misinterpretations²³. Notably, such setups can help address one of the common challenges in ceramic inspection – distinguishing true defects from structural features like grain boundaries. These boundaries, which are interfaces between crystals within the metal-ceramic substrate, often exhibit textural or reflectance characteristics that mimic surface flaws under standard

lighting conditions. By revealing complementary views of the same region under varied lighting angles, the multi-illumination strategy aids in reducing false positives associated with grain boundary misidentification. Recent advances in computer vision systems for surface defect detection have made significant progress across many fields^{7–9}. However, when computer-based inspection methods are employed, defects may be incorrectly identified as grain boundaries, which are interfaces between multiple crystals within the metal-ceramic substrate that exhibit similar characteristics²⁷.




Fig. 1: Multi-illumination setup for enhanced defect visibility in ceramic surface imaging. Combination of different lighting conditions to acquire images. R means Ring light, P-C means Pseudo-Coaxial light, F-D means Flat-Dome light²³.

Texture analysis methods play a vital role in surface inspection, particularly for ceramic materials where surface properties need to be carefully evaluated. According to the comprehensive review by Xie²⁸, statistical texture analysis approaches can be broadly categorized into several key methods that have shown promising results in defect detection. One of the most established techniques is the Gray Level Co-occurrence Matrix (GLCM), which analyzes spatial relationships between pixels at specified distances and orientations. Several studies have demonstrated its effectiveness in ceramic inspection. For example, Iivarinen *et al.*²⁹ successfully applied co-occurrence texture features to detect defects in paper web manufacturing, where the normal textures exhibited characteristic frequency patterns. The GLCM method extracts important texture features such as energy, entropy, contrast, homogeneity, and correlation, providing a rich set of descriptors for defect detection. Local Binary Patterns (LBP) represent another powerful statistical approach that has gained significant attention. LBP operates by comparing each pixel with its neighbors, creating a binary code that characterizes local texture patterns. The LBP operator has demonstrated particular strengths in being relatively invariant to illumination changes and computationally efficient³⁰. However, when applied to ceramic tile surfaces with randomly applied textures, LBP showed poorer performance compared to co-occurrence matrix and filtering-based approaches. As shown in Fig. 2, different types of surfaces present various texture challenges – from the subtle variations in ceramic tiles to more pronounced defects in materials like steel and wood. This diversity of

surface characteristics necessitates careful selection of texture analysis methods. Beyond basic statistical measures, autocorrelation features have been explored for detecting repetitive texture patterns, particularly useful in materials with regular structures. Huang and Chan³¹ demonstrated the application of fourth-order cumulants to extract harmonic peaks and successfully localize defects in textile images. However, autocorrelation functions generally prove unsuitable for random textures with irregularly arranged elements, which is often the case with ceramic surfaces. Histogram-based approaches, which though simple, have proved effective in various applications. For instance, Ng³² proposed a histogram separation technique based on Otsu's method for surface defect segmentation. However, this approach assumes that the intensity of defective regions is separable from normal regions, which isn't always true for textured surfaces. More sophisticated statistical approaches have emerged, combining multiple texture features for improved accuracy. For example, some researchers have integrated co-occurrence matrices with wavelet transform coefficients to enhance defect detection capabilities³³. This combination leverages both the spatial relationship information from GLCM and the multi-scale analysis capabilities of wavelets.

(2) Traditional machine learning approaches

The evolution of computer vision brought forth machine learning approaches that attempted to automate the feature extraction and classification process. These methods typically involve a two-stage process: feature engineering followed by classification. Feature engineering represents one of the most challenging aspects of traditional machine learning approaches. It requires domain experts to manually design and select relevant features that can effectively distinguish between defective and non-defective regions. Common features include statistical measures, geometric properties, texture descriptors, and spectral characteristics. For example, in laser-based additive manufacturing (LBAM), researchers have utilized various statistical features extracted from melt pool thermal images to detect anomalies in the printing process³⁴. The success of these systems heavily depends on the quality and discriminative power of the engineered features, making feature selection a critical but time-consuming process. In practice, feature extraction often involves multiple steps of preprocessing and transformation, such as edge detection, morphological operations, and spectrum analysis. For instance, Khanzadeh *et al.*³⁵ demonstrated this approach by extracting thermal features from melt pool images to predict porosity in direct laser deposition, achieving a prediction accuracy of 97.97 % using polynomial kernel SVM. As shown in Fig. 3, their methodology illustrates the typical workflow of traditional machine learning approaches, where thermal images undergo preprocessing and feature extraction before being fed into classification algorithms. Classification methods in traditional machine learning frameworks employ various algorithms to categorize regions based on extracted features. SVM have been widely used due to their ability to handle high-dimensional feature spaces and find optimal decision boundaries between de-

fect classes. For instance, Gobert *et al.*³⁶ applied SVM to detect defects on staircase cylinders fabricated with stainless steel GP-1 powder in the LPBF process, achieving over 80 % accuracy in identifying flaws larger than 47 μm . Other popular classifiers include k-Nearest Neighbors (k-NN), decision trees, and random forests. Garcia-Moreno *et al.*³⁷ demonstrated the effectiveness of random forest algorithms in the LMD printing process, achieving accuracies of 96.45 %, 94.56 %, and 92.21 % for macropores, micropores, and elongated pores, respectively. Traditional machine learning approaches have also shown promise in real-time monitoring applications. For example, Montazeri *et al.*³⁸ used k-NN with graph Fourier transform coefficients to achieve 90 % accuracy in layer porosity-

level prediction for nickel alloy 718 components. While these methods have demonstrated reasonable success in controlled environments where defect patterns are well-defined and consistent, their performance often deteriorates when faced with variations in lighting, texture, or defect appearance that weren't explicitly accounted for in the feature engineering process. This limitation becomes particularly evident in complex manufacturing environments where process parameters and environmental conditions can vary significantly. Additionally, the need for extensive feature engineering and domain expertise makes these approaches less adaptable to new types of defects or different manufacturing processes without substantial reconfiguration and optimization.



Fig. 2: Representative surface defects across various industrial materials. Illustrated from left to right are typical defect patterns observed on steel, stone, textile, wood, and ceramic tile surfaces. This comparison highlights the diversity of defect morphologies and textures that deep learning models must generalize across²⁸.




Fig. 3: Workflow for porosity prediction in additive manufacturing using supervised machine learning. The figure depicts the end-to-end process including data acquisition, preprocessing, feature extraction, and classification, exemplifying how statistical learning techniques have been applied in defect detection³⁵.

(3) *Limitations Driving the Need for Deep Learning*

Several fundamental limitations of traditional approaches have driven the transition toward deep learning methods. First, the reliance on manually engineered features makes traditional systems inflexible when encountering new types of defects or variations in existing ones. The feature engineering process requires significant domain expertise and must often be repeated or modified when production conditions change or new defect types emerge. Second, traditional methods typically struggle with complex textures and patterns common in ceramic surfaces, as pre-designed features may not capture all relevant characteristics for defect detection. The inability to handle varying environmental conditions presents another significant limitation. Changes in lighting, surface finish, or camera positioning can significantly impact the performance of traditional systems, requiring frequent recalibration and adjustment. Additionally, these methods often lack the capability to learn hierarchical features that could capture both fine details and broader contextual information simultaneously. This limitation becomes particularly apparent when dealing with defects that manifest at different scales or require consideration of surrounding context for accurate classification. The computational efficiency of traditional methods also becomes a bottleneck in high-speed production environments. The sequential nature of feature extraction followed by classification can introduce latency that makes real-time inspection challenging. Furthermore, the need to maintain and update multiple processing stages complicates system maintenance and optimization. These limitations, combined with the increasing availability of computational resources and labeled data, have created a compelling case for transitioning to deep learning approaches that can automatically learn relevant features and adapt to changing conditions more effectively.

III. Deep Learning Architectures for Ceramic Defect Detection

(1) *CNNs*

CNNs have fundamentally transformed ceramic surface defect detection by enabling automated feature learning from raw image data. The basic CNN architecture's effectiveness stems from its hierarchical feature extraction capabilities, which closely mirror human visual processing. In ceramic inspection, these networks have proven particularly adept at identifying a wide range of surface defects, from microscopic cracks to larger structural abnormalities. As shown in Fig. 4, CNN architectures for ceramic defect detection can be broadly categorized into two groups: traditional CNNs and custom CNNs. Traditional CNNs refer to standard, off-the-shelf architectures initially designed for general-purpose image recognition tasks – examples include LeNet, AlexNet, VGG, and ResNet. These models are often applied directly or fine-tuned for ceramic inspection without major structural modifications. In contrast, custom CNNs are specifically developed or extensively modified for the domain of ceramic surface defect detection. Such architectures typically incorporate domain-driven design elements, such as

multi-scale feature extraction modules, attention mechanisms, or specialized preprocessing pipelines, to better capture subtle defects and adapt to the complex textures characteristic of ceramic surfaces. The prevalence of custom CNNs – accounting for 67.80 % of implementations according to our survey – underscores the field's focus on tailoring models to meet the unique challenges of ceramic inspection tasks, such as high intra-class variability and limited data availability. Traditional CNN architectures typically employ a sequence of convolutional layers that progressively learn more abstract features – from basic edges and textures in initial layers to complex defect patterns in deeper layers. For instance, Min *et al.*³⁹ demonstrated the effectiveness of basic CNN architectures in ceramic defect detection using ResNet variants (ResNet20, ResNet56, and ResNet110), achieving high accuracy in identifying common defects like breaks, cracks, and pinholes. The fundamental success of these architectures lies in their ability to automatically learn relevant features without requiring manual feature engineering, making them particularly suitable for the varied nature of ceramic surface defects.

The evolution of CNN architectures in ceramic defect inspection has been marked by significant adaptations to address industry-specific challenges. Advanced variants have emerged to tackle the unique complexities of ceramic surface inspection, where defect patterns can be subtle and highly variable. For example, Karangwa *et al.*⁴⁰ presented a sophisticated approach using Faster R-CNN with VGG16 for detecting surface defects in ceramic materials, incorporating region proposal networks to improve localization accuracy. This advancement is particularly significant as it addresses one of the key challenges in ceramic inspection – the need for precise defect localization. The integration of transfer learning techniques, as evidenced by 83.05 % of studies utilizing this approach according to the paper's analysis, has become crucial in optimizing these architectures for ceramic-specific applications. The effectiveness of transfer learning is particularly notable in scenarios with limited training data, a common constraint in industrial ceramic inspection settings. Birlutiu *et al.*⁴¹ further demonstrated this trend by developing an automated defect management system with real-time high-speed processing capabilities, showcasing how modern CNN architectures can be optimized for industrial deployment while maintaining high accuracy.

Recent developments in CNN architectures have increasingly focused on addressing the specific challenges of industrial ceramic inspection environments. The incorporation of attention mechanisms and multi-scale feature fusion has proven particularly effective in handling the diverse nature of ceramic defects. These advanced architectural elements help networks focus on relevant image regions while maintaining sensitivity to defects of varying sizes and characteristics. The paper's analysis reveals that data augmentation techniques are employed in 59.32 % of studies, indicating their crucial role in enhancing model robustness. This is particularly relevant in ceramic inspection, where obtaining large, balanced datasets can be challenging. Modern CNN implementations often combine




Fig. 4: Taxonomy of CNN architectures used in ceramic surface defect detection ¹⁹.

multiple architectural innovations to achieve optimal performance. For instance, when dealing with highly reflective ceramic surfaces, specialized lighting considerations must be integrated into the model design, as demonstrated by ⁴⁰, where the authors detailed their solution using a light source intensity controller and coaxial lights to create an effective lighting system. These architectural adaptations, combined with industry-specific optimizations, have led to CNN-based systems that can reliably detect and classify ceramic surface defects in real-world manufacturing environments, significantly improving quality control processes while reducing the reliance on manual inspection.

(2) Region-based architectures

Region-based architectures have revolutionized automated visual inspection systems, particularly in industrial defect detection applications. The R-CNN family of architectures represents a major advancement by introducing region proposal mechanisms that can effectively identify and localize multiple types of surface anomalies. The original R-CNN architecture, while groundbreaking, faced computational efficiency challenges in industrial settings due to its multi-stage pipeline processing approach. Fast R-CNN addressed these limitations by enabling shared convolutional features across proposals, while Faster R-CNN made a significant leap forward by introducing the Region Proposal Network (RPN) that

generates proposals directly within the network architecture ⁴². According to Wang *et al.* ⁴³, implementing Faster R-CNN in cloud-edge computing environments achieved 72 % detection and 81 % classification accuracy for small defects in geometrically complex products such as metallic valve bodies and engine components (Fig. 5). These parts feature curved surfaces, multi-angled geometries, and high reflectivity, all of which complicate traditional visual inspection. This demonstrates the architecture's capability to handle challenging industrial inspection scenarios where defect characteristics and surface morphology can vary significantly.

The evolution of region-based architectures has been particularly important for real-time industrial applications. YOLO variants have gained prominence due to their single-shot detection approach and real-time processing capabilities. Unlike the two-stage methodology of R-CNN architectures, YOLO treats defect detection as a regression problem, directly predicting bounding boxes and class probabilities. Li *et al.* ²⁶ demonstrated that implementing improved YOLO detection networks achieved significant improvements in the speed and accuracy of steel strip surface defect detection. The architecture's ability to process multiple regions simultaneously makes it particularly valuable for high-speed production lines where real-time inspection is crucial. Recent developments have focused on enhancing the ability of these architectures to handle scale variations in defects.




Fig. 5: Prediction results using Faster R-CNN on complex surfaces. Green boxes represent predicted defect locations, while red boxes denote ground truth annotations⁴³.

For instance, Cha *et al.*⁴⁴ implemented a Faster R-CNN-based visual inspection method that achieved impressive detection rates: 90.6 %, 83.4 %, 82.1 %, 98.1 %, and 84.7 % average precision for five different defect types. These results highlight how region-based architectures can be effectively adapted for specific industrial inspection requirements. The integration of region-based architectures with traditional image processing techniques, as demonstrated by Zhou *et al.*⁴⁵, who combined saliency detection methods with template matching, achieved an 88.83 % detection rate, showing how hybrid approaches can leverage the strengths of multiple methodologies. This combination of traditional and deep learning-based approaches represents a promising direction for future developments in automated visual inspection systems.

(3) Segmentation networks

Fully Convolutional Networks (FCNs) have introduced pixel-wise prediction capabilities crucial for precise defect boundary determination in ceramic inspection. By replacing fully connected layers with convolutional layers, FCNs maintain spatial information throughout the network, enabling detailed defect segmentation. As demonstrated by Yang *et al.*⁴⁶, the FCN architecture consists of two main components: a down-sampling path containing convolutional layers, pooling layers, and dropout layers, and an up-sampling path primarily composed of deconvolutional layers. This unique end-to-end structure combines both local and global dependencies, allowing FCNs to detect defects at varying scales without slid-

ing windows across the input image. As Fig. 6 shows, the convolutional operation in FCNs processes input images through multiple channels, generating feature maps that capture increasingly complex defect patterns. The deconvolutional layers then upsample these features, reconstructing detailed spatial information critical for precise defect localization. The skip connections in FCNs help combine fine-grained details with semantic information, resulting in more accurate defect boundary delineation. In practical applications, FCNs have demonstrated remarkable efficiency – while traditional CNN-based approaches like CrackNet required nine days of training, FCN implementations achieved comparable results in less than an hour. This architecture has proven particularly valuable in applications requiring precise defect area measurements or when defect shape characteristics are critical quality indicators. The pixel-wise prediction capability enables FCNs to achieve high accuracy (97.96 %), precision (81.73 %), and recall (78.97 %) in defect detection tasks, though performance can be affected by factors such as image resolution and defect scale variations. For optimal results, the learning rate and batch size need careful tuning – studies have shown an initial learning rate of 10⁻⁴ with small batch sizes (1–4 images) provides good convergence while maintaining computational efficiency. The architecture's ability to process both grayscale and RGB images makes it versatile for different inspection scenarios, though very thin defects and complex background textures can still pose challenges requiring additional pre-processing or model refinements.

U-Net variants have emerged as powerful tools for ceramic defect segmentation, particularly in cases where precise defect boundaries are crucial. The characteristic U-shaped architecture, with its contracting and expanding paths connected by skip connections, enables effective combination of contextual information with precise localization. Modified U-Net architectures incorporating attention mechanisms or dense connections have shown enhanced performance in handling complex ceramic textures and varying defect patterns. This has been demonstrated in

recent studies focusing on ceramic crack detection, where specialized pre-processing steps combined with U-Net architectures have achieved significant improvements in segmentation accuracy. For instance, Liu *et al.*⁴⁷ proposed a WPU-Net architecture that effectively integrates spatial consistency information between adjacent slices to address challenges like blurred boundaries and missing features in ceramic surface inspection. As shown in Fig. 7, their proposed architecture incorporates a multi-level fusion strategy that allows the network to utilize both simple boundary information in the head layer and high-dimensional topology information in deeper layers, resulting in more robust crack detection. Their approach achieved an IoU of 86.5 %, significantly outperforming traditional U-Net implementations, which typically achieved around 68 % IoU on ceramic crack datasets. These variants often employ specialized loss functions designed to address class imbalance issues common in ceramic defect detection. For example, adaptive boundary weighted loss functions have been developed to better preserve geometric and topological characteristics of defects while tolerating minor differences in boundary locations⁴⁷. In another implementation focusing on ceramic surface inspection, researchers have demonstrated that U-Net based architectures can effectively handle challenges specific to ceramic materials, such as accurately detecting cracks even when they intersect with grout lines or appear in varying lighting conditions². The effectiveness of these approaches is further enhanced with data augmentation techniques and strategic model initialization methods, such as using pre-trained weights from networks like ResNet34 or ResNet50, which have shown improved generalization capabilities in ceramic defect detection tasks⁴⁸. Recent studies have also shown that when combined with appropriate pre-processing steps such as histogram equalization and Gaussian filtering, U-Net variants can achieve precision rates as high as 99.9 % in identifying ceramic surface defects, making them particularly suitable for automated quality control applications in ceramic manufacturing environments.




Fig. 6: Illustration of a convolutional layer operation in FCNs⁴⁶.




Fig. 7: Architecture and workflow of the WPU-Net model for ceramic crack detection ⁴⁷.

(4) Comparison of architectural approaches

Each architectural approach presents distinct advantages and limitations in ceramic defect detection applications. Basic CNNs offer simplicity and computational efficiency but may struggle with precise defect localization and multiple defect instances. Their strength lies in straightforward classification tasks where defect location precision is less critical. Advanced CNN variants provide improved feature learning capabilities and can handle more complex defect patterns, though at the cost of increased computational complexity and training data requirements.

Region-based architectures excel in scenarios requiring both defect localization and classification. R-CNN variants offer high detection accuracy but may face speed limitations in high-throughput inspection systems. You Only Look Once (YOLO) architectures provide faster processing speeds suitable for real-time inspection but might sacrifice some localization accuracy compared to two-stage detectors. The choice between these approaches often depends on specific industrial requirements regarding processing speed and detection accuracy.

Segmentation networks offer the most precise defect boundary detection but typically require more detailed annotation effort during training. FCNs provide good balance between segmentation accuracy and computational efficiency, while U-Net variants excel in preserving fine details crucial for certain ceramic defect types. However, these architectures may demand more computational resources and training time compared to simpler classification or detection approaches.

The selection of appropriate architecture for ceramic defect detection depends on various factors including de-

fect characteristics, inspection speed requirements, available computational resources, and annotation capabilities. Applications requiring high-speed inspection might favor YOLO variants, while those demanding precise defect boundary information might opt for segmentation networks. For example, YOLOv5s has achieved impressive real-time detection with over 96 % accuracy and 85 % mean average precision (mAP) in ceramic tile inspection tasks, making it a strong candidate for inline applications. In contrast, U-Net-based segmentation models, particularly those adapted for multi-scale feature extraction, have demonstrated superior pixel-level accuracy, with intersection over union (IoU) values exceeding 86.5 % for crack segmentation, and F1-scores above 88 % in challenging surface conditions. Fully Convolutional Networks (FCNs) also report strong performance with accuracy rates nearing 98 % and recall rates around 79 %, balancing segmentation quality with computational efficiency. Faster R-CNN, while slightly more computationally demanding, has shown over 80 % classification accuracy for small-scale defect localization tasks. These metric-driven comparisons underscore the strengths of different architectures – speed and throughput in YOLO, segmentation precision in U-Net, and balanced performance in FCNs – allowing practitioners to make architecture selections aligned with specific inspection objectives and constraints. Some implementations combine multiple architectures to leverage their complementary strengths, such as using YOLO for rapid defect detection followed by U-Net for precise defect segmentation in identified regions of interest.

IV. Implementation Considerations

The successful deployment of deep learning systems for ceramic surface defect inspection extends beyond selecting advanced network architectures. It requires a comprehensive approach that addresses various practical considerations, including data acquisition, preprocessing, model training strategies, performance evaluation, real-time processing requirements, hardware deployment, and system maintenance. Each of these aspects is critical to ensure that the developed system not only achieves high accuracy but also integrates seamlessly into existing industrial workflows.

Accurate data annotation is a foundational requirement for training supervised deep learning models in ceramic defect inspection. Given the subtle and localized nature of many ceramic defects, high-quality annotations are essential for reliable model performance. Annotation strategies can be broadly categorized into manual, semi-automatic, and synthetic methods. Manual annotation remains the most widely used approach due to its high accuracy, especially in identifying fine defects on textured ceramic surfaces. Tools such as LabelImg and CVAT developed by Intel are commonly used in ceramic defect contexts to create bounding boxes or segmentation masks. These tools provide flexible interfaces for labeling complex defect morphologies and support export in formats compatible with popular deep learning frameworks like YOLO, TensorFlow, and COCO. Semi-automatic annotation techniques are increasingly employed to reduce labeling effort. These approaches use pre-trained models or weak supervision to suggest initial labels, which are then refined manually. For instance, weakly supervised models can propose likely defect regions, while classical edge detection or thresholding can assist in delineating boundaries in low-contrast images. These methods strike a balance between efficiency and accuracy. In cases where real-world data are scarce or certain defect types are underrepresented, synthetic data generation has become an effective strategy. This includes both image augmentation and generative techniques such as GANs to produce realistic defect textures. Some studies have employed domain-specific simulators or image blending techniques to artificially embed cracks or pinholes into defect-free images, enabling robust model training in the absence of extensive labeled datasets.

Data acquisition forms the foundation of any deep learning system for defect detection. High-quality, representative data are imperative for training robust models capable of generalizing to real-world scenarios. In the context of ceramic surface inspection, capturing subtle defect characteristics necessitates careful consideration of imaging conditions. Consistent illumination is vital to minimize shadows and reflections that can obscure defects or mimic them, leading to false detections. Camera specifications, such as resolution and sensor sensitivity, must be adequate to discern fine details. For example, Wan *et al.*⁶ highlighted the challenges posed by small defect sizes and complex background textures in ceramic tiles. They employed high-resolution cameras and controlled lighting setups to ensure that defects, some averaging just 14×14 pixels,

were captured effectively without being confounded by the intricate patterns inherent to ceramic surfaces.

The preprocessing pipeline is essential to prepare the raw images for input into neural networks. This involves several steps aimed at enhancing image quality and standardizing inputs. Image normalization ensures that pixel intensity values are scaled consistently across the dataset, which is crucial for the stability of training processes. Noise reduction techniques, such as Gaussian or median filtering, mitigate the impact of sensor noise or environmental factors like dust and vibrations common in industrial settings. Color space transformations, such as converting RGB images to grayscale or other color models like HSV, can highlight features of interest. In ceramic inspection, adaptive histogram equalization may be used to enhance the contrast of defects against varying backgrounds. Preprocessing may also involve enhancing specific defect features beyond standard normalization or noise reduction. For instance, Hanzaei *et al.*⁷ applied statistical methods focused on modeling the expected radii of defects to improve edge detection performance. Their approach operated after initial preprocessing steps, using prior knowledge about typical defect geometries to better delineate defect boundaries in complex ceramic surface images.

Training deep learning models effectively requires addressing the challenges posed by industrial datasets, which often suffer from class imbalance and limited samples of defective instances. Data augmentation is a pivotal strategy to enhance dataset diversity and improve model generalization. Basic augmentation methods include geometric transformations – rotation, scaling, flipping, and translation – which simulate the variability of defect appearances due to changes in orientation or perspective. Intensity transformations adjust brightness, contrast, and saturation to mimic different lighting conditions. Advanced augmentation techniques involve generating synthetic defects or employing generative adversarial networks (GANs) to create realistic defect images, thereby addressing class imbalance. Wan *et al.*⁶ applied the Albumentations library to augment their dataset, particularly enhancing underrepresented defect classes such as white dot defects and light color block defects. This approach expanded their dataset to over 26 000 images, improving the robustness of their model.

Transfer learning has emerged as an invaluable approach when dealing with limited labeled data, which is a common scenario in industrial applications due to the high cost of manual annotation. By leveraging models pre-trained on large-scale datasets like ImageNet, one can utilize the learned feature representations that capture generic visual patterns. Fine-tuning these models on the specific domain of ceramic defects allows for adaptation to the unique characteristics of the defects. This method significantly reduces the training time and data requirements. Wan *et al.*⁶ utilized YOLOv5s – a model pre-trained on a vast dataset of general images – as their base network. They fine-tuned it with their specific dataset, implementing modifications such as deepening the backbone network and adding specialized detection layers to enhance performance on small defect targets.

Performance evaluation of the trained models must go beyond simple accuracy metrics to provide a comprehensive understanding of detection capabilities. Given the typically imbalanced nature of defect datasets – where non-defective samples vastly outnumber defective ones – metrics like precision, recall, and the F1-score become crucial. Precision measures the proportion of true positives among all positive predictions, indicating the model's ability to avoid false alarms, which is critical in reducing unnecessary inspections or discards. Recall, on the other hand, assesses the model's ability to identify all actual defects, ensuring that defective products do not proceed in the production line. The F1-score provides a balance between precision and recall. Additionally, Mean Average Precision (mAP) is essential for object detection tasks as it summarizes the precision-recall curve over all classes. Intersection over Union (IoU) metrics are used to evaluate the localization accuracy of detected defects, which is particularly important for precise defect characterization. Aligning these evaluation metrics with specific industrial requirements enables stakeholders to understand model performance in operational terms.

Real-time processing requirements pose significant challenges in industrial settings where production lines operate at high speeds, demanding rapid and accurate inspection decisions. The computational pipeline must be optimized to minimize latency. This often involves a trade-off between model complexity and inference speed. Techniques such as model pruning, quantization, and the use of lightweight architectures can reduce computational demands. Hardware and deployment considerations are crucial for the successful integration of deep learning inspection systems into production environments. The choice between edge computing and cloud-based solutions impacts factors like latency, bandwidth, and data security. Edge computing solutions, where processing is performed locally on the factory floor, offer reduced latency and are less dependent on network stability – a critical advantage for real-time applications. Hardware selection must consider the ruggedness required for industrial conditions. Devices need to withstand environmental factors such as temperature fluctuations, dust, and vibrations. Moreover, selecting energy-efficient hardware can reduce operational costs. The system must integrate seamlessly with existing production lines, requiring compatibility with industrial communication protocols and control systems⁴⁹. Interfaces for operators should be intuitive, providing clear visualizations of detected defects and allowing for easy system interaction.

The success of the implementation also depends heavily on the establishment of robust validation procedures and maintenance protocols. Regular calibration of the imaging system ensures consistency in data acquisition over time, which is vital for maintaining model performance. Performance monitoring enables the early detection of model drift, which can occur due to changes in production materials, processes, or environmental conditions. Implementing procedures for model updates – such as retraining with new data or adjusting to new types of defects – is essential for the system's adaptability. Training programs

for operational staff are necessary to ensure proper utilization and allow for quick troubleshooting. Documentation should cover system functionality, maintenance procedures, and guidelines for handling exceptions. Scalability and adaptability considerations involve designing the system architecture to accommodate future expansions, such as increased production volumes or the introduction of new product lines and defect types.

V. Challenges and Future Directions

The implementation of deep learning approaches in ceramic surface defect detection faces several interconnected technical challenges that warrant focused research attention. A primary challenge lies in the complex interaction between ceramic surface characteristics and imaging conditions, particularly in industrial settings. As evidenced by the work of Wan *et al.*⁶, even high-resolution imaging systems struggle with defects averaging 14×14 pixels, especially when these defects occur against intricate ceramic patterns. This challenge is compounded by the presence of grain boundaries, which, as noted by multiple researchers^{23,27}, can be incorrectly identified as defects due to their similar visual characteristics. The industrial environment introduces additional complications through varying lighting conditions and surface reflectance properties, which can significantly impact detection accuracy. These challenges are particularly acute in high-speed production environments where real-time processing is essential. The computational demands of sophisticated architectures like Faster R-CNN, which achieved 72 % detection and 81 % classification accuracy for small defects according to Wang *et al.*⁴³, must be balanced against the speed requirements of industrial production lines. This balance becomes even more critical when considering the need for precise defect boundary determination, as demonstrated by the WPU-Net architecture's achievement of 86.5 % IoU⁴⁷, which requires significant computational resources for real-time processing.

The data-centric challenges in ceramic defect detection present equally significant hurdles that intersect with architectural limitations. The ceramic industry's inherent class imbalance problem, where defective samples represent a minimal fraction of production output, creates fundamental difficulties in model training and validation. This challenge is evident in the implementation strategies described by various researchers, including Song and Yan²⁴, who attempted to address this with improved LBP methods combined with SVM approaches. The limited availability of labeled defect data, particularly for rare defect types, constrains the development of robust detection systems. While data augmentation techniques have shown promise, with studies reporting the use of the Albumenations library to expand datasets to over 26 000 images⁶, these approaches may not fully capture the complex variability of real-world defects. The effectiveness of transfer learning, demonstrated by the adaptation of YOLOv5s to ceramic defect detection, suggests a partial solution but raises questions about the applicability of features learned from general image datasets to the specific challenges of ceramic surface inspection. Traditional machine learning ap-

proaches achieved notable success in specific applications, such as the 97.97 % accuracy in porosity prediction using polynomial kernel SVM³⁵, but these results often relied on carefully engineered features and controlled conditions that may not generalize well to diverse production environments.

Looking toward future directions, several promising research avenues emerge from the current limitations and industry needs. The integration of physics-informed neural networks with traditional image processing techniques presents an opportunity to incorporate domain knowledge about ceramic material properties and defect formation mechanisms into detection systems – for example, by embedding thermal stress patterns or crack propagation models into the training process. This approach could help address the challenge of distinguishing between genuine defects and natural surface variations like grain boundaries, which current systems struggle to differentiate reliably. The development of adaptive architectures capable of real-time adjustment to varying environmental conditions and surface characteristics represents another crucial research direction. Such systems could potentially leverage the success of multi-illumination approaches while reducing system complexity through intelligent adaptation. The advancement of lightweight architectures optimized for edge computing deployment, building on the success of YOLO variants in real-time applications²⁶, could help address the speed-accuracy trade-off in industrial implementations. Future research should also focus on developing more sophisticated data synthesis techniques that can generate physically accurate defect patterns, potentially with the use of material-science-informed generative models that simulate defects based on realistic material failure processes such as microcrack nucleation, grain boundary fracture, or thermal fatigue. This could help address the data scarcity challenge while ensuring that synthetic data accurately represents real-world defect characteristics. The evolution of segmentation networks, particularly in handling multi-scale features and preserving fine details, as demonstrated by recent U-Net variants⁴⁷, suggests that future architectures might achieve even higher precision in defect boundary determination while maintaining computational efficiency suitable for industrial deployment.

List of Abbreviations

AI	Artificial Intelligence
CNN	Convolutional Neural Network
FCN	Fully Convolutional Network
IoU	Intersection over Union
mAP	mean Average Precision
PVD	Physical Vapor Deposition
R-CNN	Region-based Convolutional Neural Network
SVM	Support Vector Machine
U-Net	U-shaped Convolutional Network
YOLO	You Only Look Once

References

- 1 Zhao, Z.: Review of non-destructive testing methods for defect detection of ceramics, *Ceram. Int.*, **47**, 4389–4397, (2021). doi: <https://doi.org/10.1016/j.ceramint.2020.10.065>
- 2 Huang, C.-Y., Lin, I.-C., Liu, Y.-L.: Applying deep learning to construct a defect detection system for ceramic substrates, *Appl. Sci.*, **12**, 2269, (2022). doi: <https://doi.org/10.3390/app12052269>
- 3 Dong, G., Sun, S., Wang, Z., Wu, N., Huang, P., Feng, H., Pan, M.: Application of machine vision-based NDT technology in ceramic surface defect detection – a review, *Mater. Test.*, **64**, 202–219, (2022). doi: <https://doi.org/10.1515/mt-2021-2012>
- 4 Kesharaju, M., Nagarajah, R., Zhang, T., Crouch, I.: Ultrasonic sensor based defect detection and characterisation of ceramics, *Ultrasonics*, **54**, 312–317, (2014). doi: <https://doi.org/10.1016/j.ultras.2013.07.018>
- 5 Lu, Q., Lin, J., Luo, L., Zhang, Y., Zhu, W.: A supervised approach for automated surface defect detection in ceramic tile quality control, *Adv. Eng. Inform.*, **53**, 101692, (2022). doi: <https://doi.org/10.1016/j.aei.2022.101692>
- 6 Wan, G., Fang, H., Wang, D., Yan, J., Xie, B.: Ceramic tile surface defect detection based on deep learning, *Ceram. Int.*, **48**, 11085–11093, (2022). doi: <https://doi.org/10.1016/j.ceramint.2021.12.328>
- 7 Hanzaei, S.H., Afshar, A., Barazandeh, F.: Automatic detection and classification of the ceramic tiles' surface defects, *Pattern Recognit.*, **66**, 174–189, (2017). doi: <https://doi.org/10.1016/j.patcog.2016.11.021>
- 8 Chen, W., Zou, B., Huang, C., Yang, J., Li, L., Liu, J., Wang, X.: The defect detection of 3D-printed ceramic curved surface parts with low contrast based on deep learning, *Ceram. Int.*, **49**, 2881–2893, (2023). doi: <https://doi.org/10.1016/j.ceramint.2022.09.272>
- 9 Cumbajin, E., Rodrigues, N., Costa, P., Miragaia, R., Frazão, L., Costa, N., Fernández-Caballero, A., Carneiro, J., Buruberri, L.H., Pereira, A.: A real-time automated defect detection system for ceramic pieces manufacturing process based on computer vision with deep learning, *Sensors*, **24**, 232, (2024). doi: <https://doi.org/10.3390/s24010232>
- 10 Fang, T., Jafari, M.A., Danforth, S.C., Safari, A.: Signature analysis and defect detection in layered manufacturing of ceramic sensors and actuators, *Mach. Vis. Appl.*, **15**, 63–75, (2003). doi: <https://doi.org/10.1007/s00138-002-0074-1>
- 11 Sun, P., Hua, C., Ding, W., Hua, C., Liu, P., Lei, Z.: Ceramic tableware surface defect detection based on deep learning, *Eng. Appl. Artif. Intell.*, **141**, 109723, (2025). doi: <https://doi.org/10.1016/j.engappai.2024.109723>
- 12 Cao, T., Song, K., Xu, L., Feng, H., Yan, Y., Guo, J.: Balanced multi-scale target score network for ceramic tile surface defect detection, *Measurement*, **224**, 113914, (2024). doi: <https://doi.org/10.1016/j.measurement.2023.113914>
- 13 Chen, W., Zou, B., Yang, G., Zheng, Q., Lei, T., Huang, C., Liu, J., Li, L.: A real-time detection system for multiscale surface defects of 3D printed ceramic parts based on deep learning, *Ceram. Int.*, **50**, 13101–13112, (2024). doi: <https://doi.org/10.1016/j.ceramint.2024.01.220>
- 14 Zhang, H., Peng, L., Yu, S., Qu, W.: Detection of surface defects in ceramic tiles with complex texture, *IEEE Access.*, **9**, 92788–92797, (2021). doi: <https://doi.org/10.1109/ACCESS.2021.3093090>
- 15 Zhou, J., Li, H., Lu, L., Cheng, Y.: Machine vision-based surface defect detection study for ceramic 3D printing, *Machines*, **12**, 166, (2024). doi: <https://doi.org/10.3390/machines12030166>
- 16 Liu, C.H., Jeyaprkash, N., Yang, C.H.: Material characterization and defect detection of additively manufactured ce-

- ramic teeth using non-destructive techniques, *Ceram. Int.*, **47**, 7017–7031, (2021).
doi: <https://doi.org/10.1016/j.ceramint.2020.11.052>
- 17 Mariyadi, B., Fitriyani, N., Sahroni, T.R.: 2D detection model of defect on the surface of ceramic tile by an artificial neural network, *J. Phys. Conf. Ser.*, **1764**, 012176, (2021).
doi: <https://doi.org/10.1088/1742-6596/1764/1/012176>
 - 18 Lu, F., Zhang, Z., Guo, L., Chen, J., Zhu, Y., Yan, K., Zhou, X.: HFENet: A lightweight hand-crafted feature enhanced CNN for ceramic tile surface defect detection, *Int. J. Intell. Syst.*, **37**, 10670–10693, (2022).
doi: <https://doi.org/10.1002/int.22935>
 - 19 Cumbajin, E., Rodrigues, N., Costa, P., Miragaia, R., Frazão, L., Costa, N., Fernández-Caballero, A., Carneiro, J., Buruberri, L.H., Pereira, A.: A systematic review on deep learning with CNNs applied to surface defect detection, *J. Imaging*, **9**, 193, (2023). doi: <https://doi.org/10.3390/jimaging9100193>
 - 20 Stephen, O., Maduh, U.J., Sain, M.: A machine learning method for detection of surface defects on ceramic tiles using convolutional neural networks, *Electron.*, **11**, 55, (2022).
doi: <https://doi.org/10.3390/electronics11010055>
 - 21 Huang, Y.P., Su, C.M., Basanta, H., Tsai, Y.L.: Imbalance modelling for defect detection in ceramic substrate by using convolutional neural network, *Processes*, **9**, 1678, (2021).
doi: <https://doi.org/10.3390/pr9091678>
 - 22 Nogay, H.S., Akinci, T.C., Yilmaz, M.: Detection of invisible cracks in ceramic materials using pre-trained deep convolutional neural network, *Neural Comput. Appl.*, **34**, 1423–1432, (2022). doi: <https://doi.org/10.1007/s00521-021-06652-w>
 - 23 Zou, M., Matsunaga, K., Ueda, Y., Sugawara, T., Osanai, H., Kageyama, Y.: Defect detection in metal-ceramic substrate based on image processing and machine learning, *IEEE J. Ind. Appl.*, 23006878, (2024).
 - 24 Song, K., Yan, Y.: A noise robust method based on completed local binary patterns for hot-rolled steel strip surface defects, *Appl. Surf. Sci.*, **285**, 858–864, (2013).
doi: <https://doi.org/10.1016/j.apsusc.2013.09.002>
 - 25 Ghorai, S., Mukherjee, A., Gangadaran, M., Dutta, P.K.: Automatic defect detection on hot-rolled flat steel products, *IEEE Trans. Instrum. Meas.*, **62**, 612–621, (2012).
 - 26 Li, J., Su, Z., Geng, J., Yin, Y.: Real-time detection of steel strip surface defects based on improved YOLO detection network, *IFAC-Pap. OnLine*, **51**, 76–81, (2018).
doi: <https://doi.org/10.1016/j.ifacol.2018.09.412>
 - 27 Ren, Z., Fang, F., Yan, N., Wu, Y.: State of the art in defect detection based on machine vision, *Int. J. Precis. Eng. Manuf. Green Technol.*, **9**, 661–691, (2022).
doi: <https://doi.org/10.1007/s40684-021-00343-6>
 - 28 Xie, X.: A review of recent advances in surface defect detection using texture analysis techniques, *ELCVIA Electron. Lett. Comput. Vis. Image Anal.*, 1–22 (2008).
 - 29 Iivarinen, J., Rauhamaa, J., Visa, A.: Unsupervised segmentation of surface defects. In: *Proc. 13th Int. Conf. Pattern Recognit.*, **4**, 356–360 (1996).
doi: <https://doi.org/10.1109/ICPR.1996.547445>
 - 30 Mäenpää, T., Pietikäinen, M.: Texture analysis with local binary patterns. In: *Handbook of Pattern Recognition and Computer Vision*. World Scientific, 197–216, (2005).
 - 31 Huang, Y., Chan, K.L.: Texture decomposition by harmonics extraction from higher order statistics, *IEEE Trans. Image Process.*, **13**, 1–14, (2004).
doi: <https://doi.org/10.1109/TIP.2003.819432>
 - 32 Ng, H.F.: Automatic thresholding for defect detection, *Pattern Recognit. Lett.*, **27**, 1644–1649, (2006).
doi: <https://doi.org/10.1016/j.patrec.2006.03.009>
 - 33 Latif-Amet, A., Ertüzün, A., Erçil, A.: An efficient method for texture defect detection: sub-band domain co-occurrence matrices, *Image Vis. Comput.*, **18**, 543–553, (2000).
doi: [https://doi.org/10.1016/S0262-8856\(99\)00062-1](https://doi.org/10.1016/S0262-8856(99)00062-1)
 - 34 Mahmoudi, M., Ezzat, A.A., Elwany, A.: Layerwise anomaly detection in laser powder-bed fusion metal additive manufacturing, *J. Manuf. Sci. Eng.*, **141**, 031002, (2019).
 - 35 Khanzadeh, M., Chowdhury, S., Marufuzzaman, M., Tschopp, M.A., Bian, L.: Porosity prediction: supervised-learning of thermal history for direct laser deposition, *J. Manuf. Syst.*, **47**, 69–82, (2018). doi: <https://doi.org/10.1016/j.jmsy.2018.04.001>
 - 36 Gobert, C., Reutzel, E.W., Petrich, J., Nassar, A.R., Phoha, S.: Application of supervised machine learning for defect detection during metallic powder bed fusion additive manufacturing using high resolution imaging, *Addit. Manuf.*, **21**, 517–528, (2018). doi: <https://doi.org/10.1016/j.addma.2018.04.005>
 - 37 García-Moreno, A.-I., Alvarado-Orozco, J.-M., Ibarra-Medina, J., Martínez-Franco, E.: Image-based porosity classification in Al-alloys by laser metal deposition using random forests, *Int. J. Adv. Manuf. Technol.*, **110**, 2827–2845, (2020).
doi: <https://doi.org/10.1007/s00170-020-05887-6>
 - 38 Montazeri, M., Nassar, A.R., Dunbar, A.J., Rao, P.: In-process monitoring of porosity in additive manufacturing using optical emission spectroscopy, *IIEE Trans.*, **52**, 500–515, (2020).
 - 39 Min, B., Tin, H., Nasridinov, A., Yoo, K.-H.: Abnormal detection and classification in i-ceramic images. In: 2020 IEEE Int. Conf. Big Data Smart Comput. (BigComp), 17–18, (2020).
doi: <https://doi.org/10.1109/BigComp48618.2020.0-106>
 - 40 Karangwa, J., Kong, L., You, T., Zheng, J.: Automated surface defects detection on mirrorlike materials by using Faster R-CNN. In: 2020 7th Int. Conf. Inf. Sci. Control Eng. (ICISCE), 2288–2294, (2020).
doi: <https://doi.org/10.1109/ICISCE50968.2020.00341>
 - 41 Birlutiu, A., Burlacu, A., Kadar, M., Onita, D.: Defect Detection in Porcelain Industry Based on Deep Learning Techniques. 2017 19th International Symposium on Symbolic and Numeric Algorithms for Scientific Computing (SYNASC), 263–270, (2017).
doi: <https://doi.org/10.1109/SYNASC.2017.00049>
 - 42 Ren, S., He, K., Girshick, R., Sun, J.: Faster R-CNN: towards real-time object detection with region proposal networks, *IEEE Trans. Pattern Anal. Mach. Intell.*, **39**, 1137–1149, (2017).
doi: <https://doi.org/10.1109/TPAMI.2016.2577031>
 - 43 Wang, Y., Liu, M., Zheng, P., Yang, H., Zou, J.: A smart surface inspection system using faster R-CNN in cloud-edge computing environment, *Adv. Eng. Inform.*, **43**, 101037, (2020).
doi: <https://doi.org/10.1016/j.aei.2020.101037>
 - 44 Cha, Y.J., Choi, W., Suh, G., Mahmoudkhani, S., Büyükköztürk, O.: Autonomous structural visual inspection using region-based deep learning for detecting multiple damage types, *Comput.-Aided Civ. Infrastruct. Eng.*, **33**, 731–747, (2018).
doi: <https://doi.org/10.1111/mice.12334>
 - 45 Zhou, X., Wang, Y., Xiao, C., Zhu, Q., Lu, X., Zhang, H., Ge, J., Zhao, H.: Automated visual inspection of glass bottle bottom with saliency detection and template matching, *IEEE Trans. Instrum. Meas.*, **68**, 4253–4267, (2019).
doi: <https://doi.org/10.1109/TIM.2018.2886977>
 - 46 Yang, X., Li, H., Yu, Y., Luo, X., Huang, T., Yang, X.: Automatic pixel-level crack detection and measurement using fully convolutional network, *Comput.-Aided Civ. Infrastruct. Eng.*, **33**, 1090–1109, (2018).
doi: <https://doi.org/10.1111/mice.12412>
 - 47 Liu, W., Chen, J., Liu, C., Ban, X., Ma, B., Wang, H., Xue, W., Guo, Y.: Boundary learning by using weighted propagation in convolution network, *J. Comput. Sci.*, **62**, 101709, (2022).
doi: <https://doi.org/10.1016/j.jocs.2022.101709>

⁴⁸ Junior, G.S., Ferreira, J., Millán-Arias, C., Daniel, R., Junior, A.C., Fernandes, B.J.T.: Ceramic cracks segmentation with deep learning, *Appl. Sci.*, **11**, 6017, (2021).
doi: <https://doi.org/10.3390/app11136017>

⁴⁹ Howard, A.G.: Mobilenets: Efficient convolutional neural networks for mobile vision applications, arXiv Prepr. arXiv, 1704.04861, (2017).