RT&A, No 4 (89)
Volume 20, December 2025

QUALITY CONTROL FOR SUSTAINABILITY IN CASTING MANUFACTURING

Manzoor A. Khanday[1*], Trushal Hirani[2], Dilawar A. Bhat[3], Vaddi Tejasvi[4], Anusha Ale Saji[5]

•

[1*,2,4,5] Department of Statistics, Lovely Professional University, Punjab, India
[3]Symbiosis School of Banking and Finance, Symbiosis International University, Pune, India
manzoorstat@gmail.com, hiranitrushal@yahoo.com, dilawar.bhat@ssbf.edu.in, vadditejasvi@gmail.com, anushasaji1002@gmail.com.
Corresponding author[1*]: manzoorstat@gmail.com

Abstract

This study proposes an artificial intelligence–driven framework for automating defect detection in casting manufacturing, with a particular focus on enhancing sustainability in the production of submersible pump impellers. A dataset comprising more than 7,000 grayscale images—classified as defective or acceptable—was utilized to train and evaluate three deep learning models: a custom-built Convolutional Neural Network (CNN), VGG16, and ResNet18. The models were assessed for their accuracy, precision, and recall, with VGG16 and ResNet18 demonstrating superior performance, achieving over 99% classification accuracy. Advanced image preprocessing techniques, including histogram equalization and Canny edge detection, were employed to enhance defect visibility and enable the detection of microstructural flaws such as blow holes, shrinkage cavities, and surface irregularities. The integration of transfer learning significantly reduced training time while improving generalization across unseen samples. From a sustainability perspective, the adoption of AI-based automated inspection reduces material wastage, rework, and energy consumption by ensuring early identification of defective products. Furthermore, the deployment of these models on edge devices and integration with IoT-enabled monitoring systems facilitate real-time, continuous quality assurance without human intervention. Overall, the findings underscore that AI-driven quality control systems not only optimize manufacturing efficiency and product reliability but also contribute to achieving broader sustainability goals by minimizing resource depletion, enhancing energy efficiency, and supporting circular manufacturing practices.

Keywords: Quality Control, Casting Manufacturing, Deep Learning, Convolutional Neural Networks, VGG16, ResNet18, Defect Detection, Edge AI, IoT Integration, Sustainability.

I. Introduction

In today's manufacturing landscape, sustainability and quality are two interlinked pillars that drive industry innovation. Traditional casting processes, though pivotal in producing complex components like submersible pump impellers, are often marred by imperfections that can lead to

DOI: https://doi.org/10.24412/1932-2321-2025-489-1189-1200

significant material waste and operational inefficiencies. Traditional casting processes are foundational to manufacturing complex components like submersible pump impellers, but they are inherently prone to defects such as blow holes, pinholes, and shrinkage anomalies. These imperfections lead to material waste, energy inefficiency, and increased production costs, undermining sustainability goals in manufacturing. Swift and Booker highlight that up to 20% of casting products are discarded due to defects, emphasizing the urgent need for improved quality control systems to minimize resource waste [1]. Manual inspection, the conventional approach for defect detection, exacerbates these challenges due to its labor-intensive nature, subjectivity, and high error rates. Tao notes that manual methods fail to scale with modern production demands, often leading to inconsistent quality assessments and delayed interventions [2]. Casting defects such as blow holes, pinholes, shrinkage anomalies, and other irregularities—pose not only a threat to product quality but also to the sustainable utilization of resources. Manual inspection, the conventional method for detecting these flaws, is labor-intensive, time-consuming, and prone to human error [3]. This inefficiency can result in the inadvertent rejection of otherwise salvageable products, thereby escalating production costs and environmental footprint. This chapter explores how the integration of deep learning techniques into quality control processes can revolutionize casting manufacturing. By leveraging a meticulously curated dataset comprising over 7,300 grayscale images of casting products—augmented to enhance the diversity of defect presentations—we demonstrate the potential for automated, real-time defect detection. The study not only evaluates defect detection accuracy but also considers its implications for sustainable production, highlighting how automated inspection contributes to resource circularity and lower carbon intensity in manufacturing lines. The dataset is divided into two principal categories: defective and acceptable castings, providing a robust framework for training and validating classification models. The automated system aims practices by reducing material waste, minimizing energy consumption, and optimizing overall production workflows. Furthermore, by transitioning from manual to machine-driven inspection, manufacturers can achieve a higher degree of precision and consistency in quality control. This shift supports a circular economy, where resources are utilized more efficiently, and waste is significantly reduced [4]. In the broader context of sustainable industry practices, the use of advanced artificial intelligence (AI) systems for quality control embodies a forward-thinking approach—melding technological innovation with environmental stewardship. In the following sections, we delve into the specifics of the dataset, outline the deep learning methodologies employed, and evaluate the performance of our classification models. Ultimately, this chapter underscores how AI-driven quality control not only enhances production efficacy but also serves as a critical enabler of sustainability in modern manufacturing.

II. Review of Literature

Recent advancements in artificial intelligence, particularly the emergence of Vision Transformers (ViTs), have significantly enhanced the ability to detect complex surface defects in manufacturing. Dosovitskiy et al. demonstrated that ViTs outperform traditional Convolutional Neural Networks (CNNs) in capturing long-range dependencies within high-resolution images, thereby improving the precision of defect localization and classification. However, their high computational demands continue to limit widespread deployment in real-time industrial applications. These challenges underscore a persistent gap between conventional quality control practices and the sustainability imperatives of Industry 4.0. Although deep learning has revolutionized manufacturing inspection systems, the integration of sustainability indicators—such as reductions in material waste, rework, and energy consumption—into AI-driven frameworks remains limited. Addressing this gap, the present research aligns deep learning performance metrics with sustainability outcomes in casting operations. Models such as CNN, ResNet, and VGG demonstrate strong potential for automated defect detection with minimal human intervention. By minimizing scrap rates and optimizing energy usage, AI-based quality control systems not only improve production efficiency but also contribute to sustainable manufacturing [5]. Bressanelli et al. emphasize that such automation plays a pivotal role in advancing circular economy principles by conserving resources and reducing environmental impact [6].

III. Sustainability Implications of AI-Based Quality Control

A key motivation for integrating deep learning into casting defect detection is the sustainability benefits it brings to manufacturing processes. First, reducing resource wastage emerges as a direct outcome of early defect identification [7]. By identifying flawed products immediately after casting, manufacturers avoid investing additional energy, labor, and materials into parts that would ultimately be scrapped, thereby lowering raw material consumption [8]. In parallel, energy efficiency gains arise from automated inspections operating in real-time, allowing production lines to run continuously without the bottlenecks of manual checks [9]. This streamlined workflow not only cuts down idle machine time but also diminishes the likelihood of large-scale rework cycles, further conserving energy across the production chain [10]. From a cost-benefit perspective, AI-driven inspection systems deliver high accuracy and consistency, minimizing false rejections of otherwise viable parts [11]. Although initial investments in hardware (e.g., GPUs) and data infrastructure can be significant, the long-term savings from reduced scrap rates and improved throughput frequently offset these costs. Collectively, these factors illustrate how advanced quality control solutions not only enhance product integrity but also support broader environmental and economic objectives within casting and other manufacturing domains.

IV. Methods

A critical foundation for automating quality control in casting manufacturing lies in the quality and structure of the dataset. The dataset employed in this study is meticulously curated to capture the nuances of casting defects that commonly occur in the production of submersible pump impellers. This section delves into the dataset's composition, acquisition process, and relevance to sustainable manufacturing practices. All images in the dataset are top-view shots of submersible pump impellers, captured under controlled lighting conditions. The dataset is organized into two primary subsets, each serving a specific purpose in the development and validation of the deep learning model.

The dataset used in this study consists of two primary subsets designed to ensure both diversity and precision in defect detection. The augmented image set (300×300 pixels) comprises 7,348 grayscale images that have undergone various augmentation techniques such as rotation, flipping, and contrast adjustment to simulate diverse real-world conditions [12], this process enriches the dataset, allowing the deep learning models to generalize effectively to unseen samples and reducing overfitting. In contrast, the original image set (512×512 pixels) contains high-resolution, unaltered images that capture intricate details of the casting process. This subset includes 519 images labelled as "OK" and 781 labelled as "Defective," offering a clear visual distinction between flawless and flawed impellers. Together, these datasets provide a comprehensive foundation for developing robust and accurate models capable of detecting subtle surface irregularities critical for automated quality control in casting manufacturing. This dual-dataset approach ensures model reliability, scalability, and adaptability for real-world industrial applications.

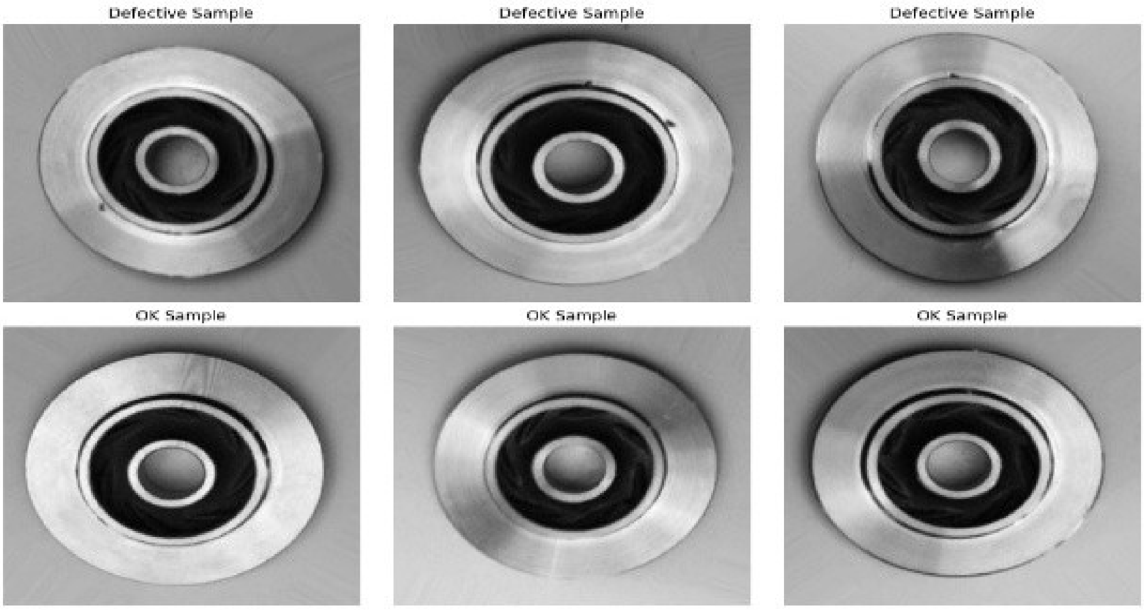

Figure 1: *Sample Images of OK and Defectives. The visible tonal variations between OK and defective impellers mark the first indication of quality differences.*

Figure 1 illustrates representative samples from both OK and defective categories. The visible tonal and textural contrasts reveal that defective impellers often display uneven surface shading, micro-pits, and darker contours compared with the smoother, uniformly lit OK counterparts. These differences form the visual cues on which the deep-learning models base their classification.

Understanding the dataset involves more than just reading numbers and images; it also reflects the broader sustainability goals of modern manufacturing. Detecting defects at an early stage helps cut down waste by identifying faulty products before they move further in the production line. This saves both raw materials and energy, which is especially valuable in casting processes where even small flaws can lead to large production losses. Automated defect detection systems that use deep learning improve energy efficiency by working faster and more reliably than manual inspection. They lower energy use and operating costs while keeping quality standards consistent. By reducing human error, limiting scrap, and minimizing rework, such systems make better use of available resources and support a circular economy that ties directly into the sustainability objectives of the manufacturing sector.

To gain an overview of the grayscale intensity values, the mean and standard deviation were computed for a subset of 100 images in each class, a few of which are presented in Figure 1. For the Defective class, the mean grayscale intensity was 138.66 with a standard deviation of 58.02, whereas for the OK class, the mean was 149.78 with a standard deviation of 62.42.

These values highlight two important patterns. The higher mean intensity in the OK class indicates that these images are generally brighter than the defective ones. This difference may arise from variations in surface reflectivity or the presence of darker regions such as pits, cracks, or holes in defective components. Both classes also show relatively high standard deviations, reflecting considerable variability in brightness and contrast. Such variability is crucial to recognize for model training, as it emphasizes the importance of applying suitable pre-processing techniques such as normalization or histogram equalization to ensure the model handles differences in illumination and texture effectively. The histograms of pixel intensities often display multiple peaks, suggesting the presence of distinct surface features or reflective highlights. In defective images, these peaks tend to be more widely spread, which may correspond to the coexistence of darker defect areas and brighter surrounding regions. Defective samples typically have a larger portion of pixel values concentrated in the lower-to-mid intensity range, consistent with their lower mean intensity of 138.66. In contrast, OK samples show histogram peaks shifted toward higher intensity values, aligning with the higher mean of 149.78. This pattern indicates that non-defective parts generally possess smoother and more uniformly bright surfaces. Even among images belonging to the same class, the brightness distribution can vary considerably. This within-class variability may result from subtle differences in lighting during image capture, minor fluctuations in the manufacturing process, or unique visual signatures of specific defects. Understanding this variability is vital for developing robust models capable of distinguishing genuine defects from variations caused by external factors.

Figures 2 and 3 illustrate the pixel intensity distributions for both defective and OK impeller images, offering key insights into their visual characteristics. The defective samples exhibit a broader and multi-modal intensity curve, reflecting uneven surface textures, cavities, and rough patches that cause irregular light reflectance. In contrast, the OK samples display a narrower, right-shifted histogram, signifying consistent brightness, smooth surface finish, and structural uniformity. These clear distinctions between the two classes emphasize the need for robust preprocessing techniques—specifically normalization and histogram equalization—to standardize illumination conditions and enhance contrast before model training and feature extraction in subsequent stages.

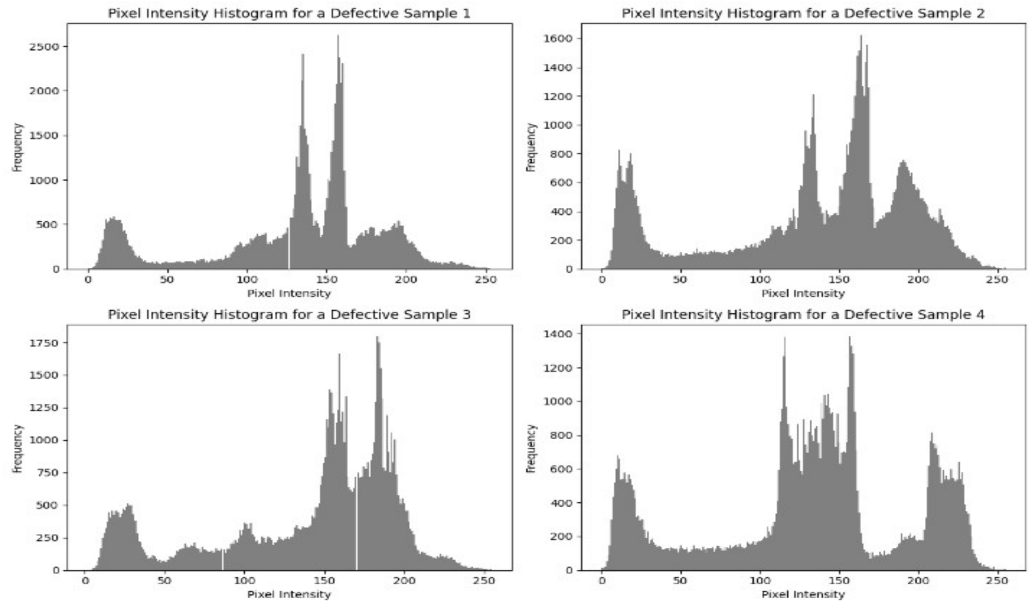

Figure 2: *Pixel Intensity for Defective Samples.*

The histogram in Figure 2 depicts the grayscale pixel intensity distribution for defective impeller images, showing a broad and irregular profile with multiple secondary peaks. This pattern indicates a wide variation in brightness values across the defective surfaces, reflecting the heterogeneous nature of casting flaws such as blow holes, shrinkage cavities, and uneven metal flow. These imperfections create localized dark and bright regions, leading to complex tonal variations. The multi-peaked structure further suggests that the defects differ in both shape and depth, adding to the overall irregularity of the images. Such variability introduces significant intra-class differences within the defective category, emphasizing the importance of image preprocessing. Techniques like normalization or histogram equalization are therefore essential to minimize intensity inconsistencies and ensure that the model focuses on meaningful defect features during training rather than illumination disparities.

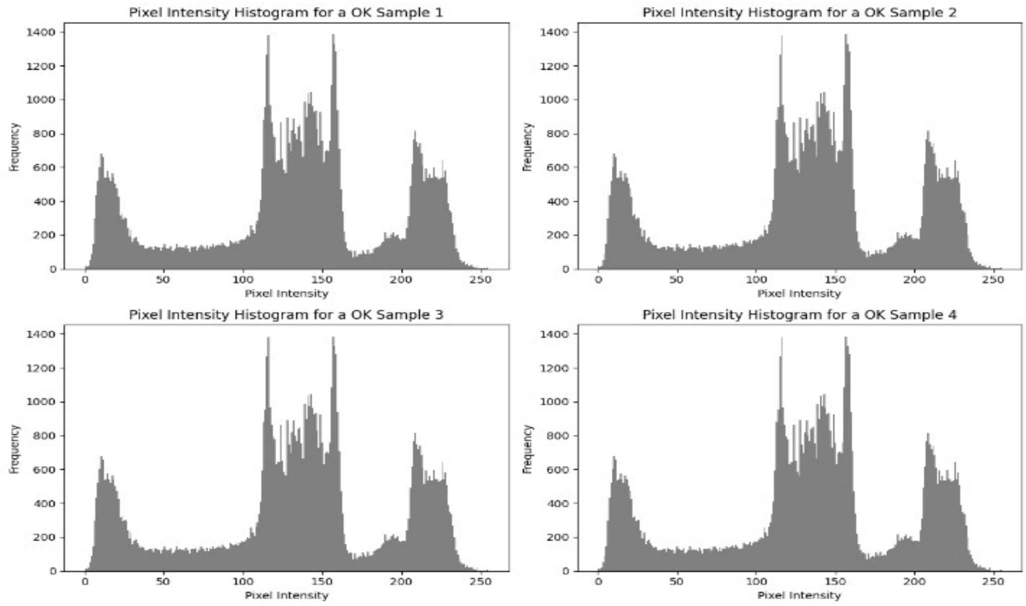

Figure 3: *Pixel Intensity for OK Samples*

Figure 3 illustrates the grayscale intensity distribution for OK impeller images, revealing a distinct pattern compared with the defective class. The histogram is notably more compact and right-shifted toward higher intensity values, signifying generally brighter, smoother, and more uniform surfaces. The presence of a single dominant peak indicates a stable and consistent reflectance pattern, which is characteristic of defect-free metallic surfaces. This narrower spread of pixel values also suggests reduced variability within the OK class, reflecting greater structural integrity and surface uniformity. Such consistent brightness behavior provides a strong separability cue for classification models, allowing them to effectively differentiate between OK and defective samples based on stable intensity features.

From a quality control standpoint, these histogram characteristics reinforce the quantitative distinctions observed in mean and standard deviation values between the two classes. Deep learning models can leverage these differences as discriminative features; however, preprocessing steps—such as normalization or histogram equalization—are essential to mitigate illumination inconsistencies and ensure that learning focuses on defect-related textures rather than simple brightness variations.

Figures 4 and 5 display edge maps obtained via the Canny operator. Continuous, symmetric rings characterize defect-free impellers, while interruptions and asymmetries denote potential casting flaws such as blow holes or burrs. These annotated edges supply geometric evidence at complements intensity-based analysis, helping the network learn both texture and shape-related defect signatures. In other words, further enhance this understanding by displaying edge maps generated using the Canny edge detection operator. In defect-free impellers, the edges appear continuous, symmetric, and well-defined, reflecting the ideal geometry of properly cast components. In contrast, defective samples exhibit irregular, fragmented, or discontinuous edges, indicating casting flaws such as blow holes, burrs, or surface distortions. These geometric variations complement intensity-based analysis, enabling the deep learning model to integrate both textural and shape-based cues for more reliable classification. Consequently, combining histogram-based brightness patterns with edge-based geometric features strengthens the robustness and interpretability of the automated defect detection framework, advancing both accuracy and sustainability in quality control operations.

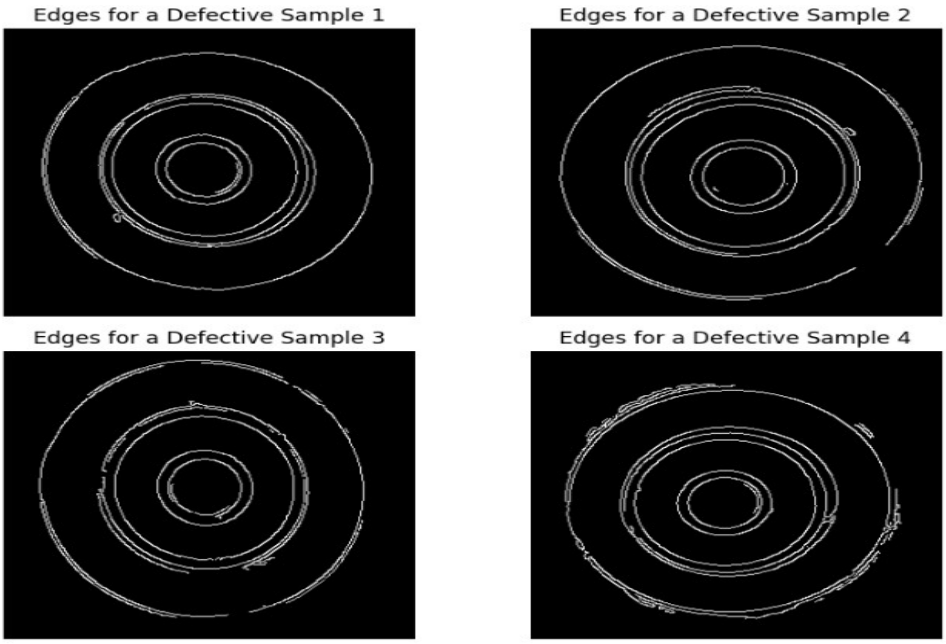

Figure 4: *Edge Detection for Defective Samples*

Figure 4 presents the Canny edge detection output for defective impeller images. The edge map reveals fragmented, uneven contours and discontinuous circular patterns across the impeller surface. These irregularities correspond to structural disruptions such as blow holes, burrs, or minor cracks formed during casting. The varying edge thickness and broken ring formations indicate deviations from ideal symmetry, a key geometric marker of surface defects. Additional localized edges appear around defect regions, signifying abrupt intensity transitions caused by rough textures or cavities. Such patterns highlight the geometric complexity of defective parts and emphasize the model's need to capture both intensity-based and shape-based cues for accurate classification.

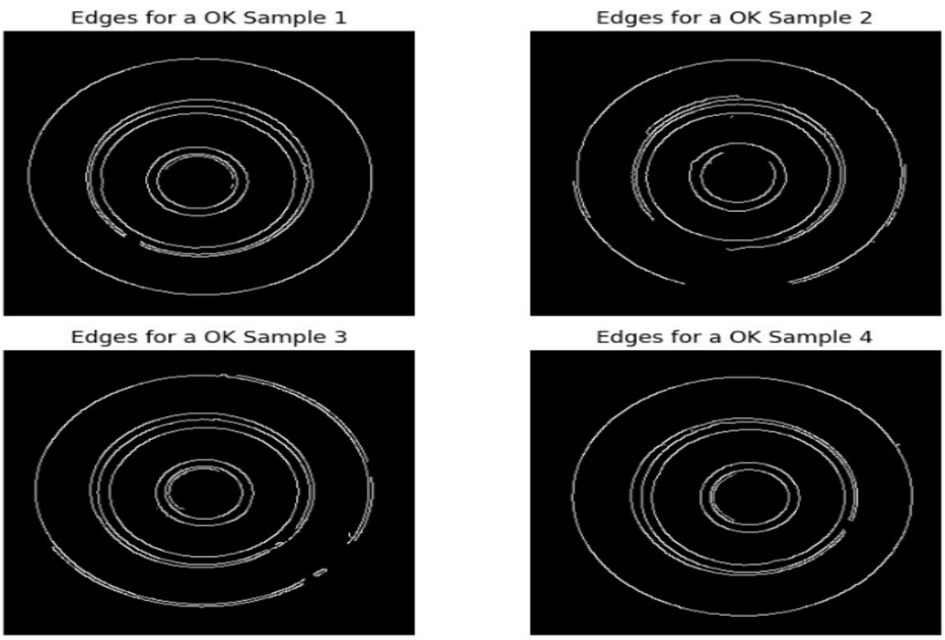

Figure 5: *Edge Detection for OK samples*

Figure 5 displays the Canny edge detection results for OK impeller images. The edges are smooth, continuous, and symmetrically arranged in concentric rings that align with the expected geometry of a properly cast impeller. The uniform edge intensity and consistent ring structure suggest minimal surface distortion and a well-balanced manufacturing outcome. The absence of stray or fragmented edges reflects the homogeneity of surface finish and stable illumination during image capture. These characteristics establish a geometric baseline against which defective samples can be contrasted. The consistency observed in the OK edges demonstrates the reliability of the imaging setup and supports the use of edge features as a dependable input for automated quality inspection.

The Canny edge detection results offer a detailed view of the geometric structure and surface quality of both defective and OK samples. Each impeller exhibits circular or near-circular regions that the Canny operator captures as concentric rings. In the OK samples, these rings appear continuous and sharply defined, indicating a smooth and consistent surface. In contrast, defective samples show subtle breaks, irregular thickness, or extra edge formations, suggesting disruptions on the surface. The Canny Edge Detection Analysis reveals the following:

- Small discontinuities, bumps, or distortions in the circular outlines often correspond to physical surface defects such as blow holes or burrs. Some defective samples display additional fragmented edges that are absent in the OK set, indicating localized imperfections or inconsistencies in material composition.
- Submersible pump impellers are designed for high symmetry, which is clearly reflected in the well-balanced edges of OK parts. Defective samples, however, sometimes deviate from this pattern, showing asymmetric or distorted contours that become more prominent through Canny edge outlines.
- Since Canny edge detection depends on intensity gradients, differences in illumination can create artificial edge patterns. Although the controlled lighting minimizes glare and shadow artifacts, residual highlights may still appear as edges in both defective and OK samples. These are consistent across images and should not be mistaken for actual defects. Edge-detected images can serve as a basis for quantitative feature extraction. Metrics such as perimeter, circularity, or the number of closed contours can be derived and integrated into a machine learning pipeline for defect classification. Consistency in edge capture across samples also indicates a stable imaging setup. Significant disparities in edge quality may signal a need to recalibrate the lighting or camera parameters.
- Detecting early signs of surface defects through edge-based analysis enables timely intervention during production. This reduces rework, material waste, and product rejection rates. Enhanced defect detection contributes to sustainable manufacturing by conserving resources and minimizing environmental impact.

The objective of this research is to build a robust automated quality inspection system for detecting defects in submersible pump impellers. We approach this task as a binary classification problem—each image is classified as either defective or OK. To accomplish this, we compare four CNN models:

- A CNN built from scratch (custom architecture)
- ResNet18 (transfer learning)
- VGG16 (transfer learning)

The CNN model we developed utilizes a sequential, layer-by-layer design tailored for binary classification (defective vs. OK). We begin with an input layer shaped to accommodate single-channel (grayscale) images, defined by input shape = (height, width, 1). The first trainable component is a convolutional layer with 32 filters, each sized 3×3, using the ReLU activation function to introduce nonlinearity. We employ stride=2 and padding='same', which effectively halves the spatial dimensions while preserving edge information. Immediately following is a max pooling layer (MaxPooling2D), which further reduces the feature map's width and height by a factor of two, helping the network learn translation-invariant features while lowering computational cost.

This pattern repeats in the second convolution block, where the number of filters increases to 64, signifying a deeper, more discriminative representation of the image [13]. Again, stride and max pooling are used to down sample feature maps efficiently. After these blocks, the Flatten layer transforms the two-dimensional feature maps into a one-dimensional vector, bridging the gap between convolutional feature extraction and fully connected classification. Next, a Dense layer with 128 neurons captures high-level interactions among the extracted features, using ReLU to maintain nonlinearity [14]. Finally, a single-output Dense layer with a sigmoid activation translates the network's learned representations into a probability for one of the two classes (defective or OK). To train this model, we compile it with the Adam optimizer at a default learning rate of 0.001, striking a balance between speed of convergence and stability. The binary cross-entropy loss function measures the divergence between predicted probabilities and actual labels, guiding backpropagation. We track accuracy as a metric for quick performance monitoring. By stacking these layer-two convolution-pooling blocks followed by flattening and fully connected layers—we achieve a compact yet powerful architecture capable of learning meaningful patterns from grayscale images. This design effectively captures local features through convolution and pooling while leveraging a dense layer to consolidate information for the final binary decision.

In parallel to our custom CNN, we employ transfer learning with two popular architectures: ResNet18 and VGG16. Each of these models is pre-trained on a large dataset (typically ImageNet), which provides a strong initial feature representation. We then fine-tune the models on our impeller images to adapt these representations to the casting defect domain. By fine-tuning them on our specific impeller image dataset, we enable the networks to adapt their learned representations to effectively identify casting defects, thereby enhancing accuracy, reducing training time, and improving overall model generalization within the defect detection domain.

I. ResNet18

ResNet18 (Residual Network with 18 layers), tackles the vanishing gradient problem in deeper networks by introducing residual learning. Unlike traditional CNNs, ResNet18 features skip connections—or shortcuts—that directly feed the input of a residual block to its output [15]. Each residual block typically includes two 3×3 convolution layers with batch normalization and ReLU activations, plus a skip connection that bypasses these layers, enabling gradients to flow more easily during backpropagation [16]. This design allows the network to learn identity mappings, reducing the difficulty of training as depth increases. Structurally, ResNet18 begins with a 7x7 convolution and max pooling, followed by four stages of residual blocks, each doubling the number of filters while halving spatial dimensions at stage boundaries. The final stage is a global average pooling layer and a fully connected layer for classification. Despite having 18 layers, ResNet18 is relatively lightweight and easier to train compared to deeper variants (e.g., ResNet50), making it well-suited for tasks where computational resources are limited but robust feature extraction is essential. These architectures, while both CNN-based, differ in their approach to depth and parameterization. VGG16 relies on a strictly layered design with small, uniform filters and deep fully connected layers, whereas ResNet18 uses skip connections to maintain gradient flow in a shallower (but more efficient) framework. Their complementary strengths—VGG16's simplicity and high capacity versus ResNet18's ease of optimization and residual learning—make them popular choices for transfer learning, fine-tuning, and real-world image classification problems [17].

For our two-class problem, we replace the final fully connected layer with one that has 2 output neurons. We also selectively unfreeze later layers for fine-tuning while keeping earlier layers fixed (or slowly updated), this design ensures that deeper layers continue learning meaningful features rather than merely memorizing patterns.

II. VGG16

VGG16, introduced by Simonyan and Zisserman (2015), is a seminal deep convolutional neural network (CNN) renowned for its simplicity, uniformity, and strong generalization capabilities. The architecture comprises 16 weight layers—13 convolutional and 3 fully connected layers—employing small, uniform 3×3 convolution filters throughout. Each convolutional block typically contains two or three convolutional layers with a stride of 1 and padding of 1 to preserve spatial resolution, followed by a max-pooling layer that reduces the feature map size by half. This hierarchical structure allows the model to progressively learn from low-level edges and textures to high-level object patterns. After feature extraction, the flattened feature maps pass through three dense layers, the first two with 4,096 neurons each and the final one producing class probabilities through a SoftMax activation. VGG16’s proven efficiency on large-scale datasets such as ImageNet makes it highly effective for transfer learning. In this study, VGG16 was used as a feature extractor, and its final layers were retrained for binary classification of submersible pump impeller defects. The model employed cross-entropy loss and the Adam optimizer with an initial learning rate of 0.001, decaying gradually through a step or exponential schedule, and was trained for 10–30 epochs with a batch size of 32 to prevent overfitting and ensure rapid convergence [14].

From the results, we observe that VGG and ResNet outperform CNN in terms of precision and recall, especially for the OK class. CNN tends to misclassify OK parts more than VGG and ResNet, as reflected in its lower precision for the OK class (0.96). VGG and ResNet have near-identical performance, both reaching 99% accuracy with minimal misclassifications. ResNet offers better generalization over unseen datasets due to its deeper architecture and residual learning.

V. Results

Table 1: *Model Performance*

Model	Accuracy	Precision (OK)	Precision (Defective)	Recall (OK)	Recall (Defective)
CNN	98%	0.96	0.99	0.98	0.98
VGG	99%	1.00	0.98	0.99	1.00
ResNet	99%	1.00	0.98	0.99	1.00

The training and validation loss and accuracy trends for CNN, VGG, and ResNet were analysed over 20 epochs. The plots reveal that CNN shows a slightly higher validation loss than VGG and ResNet, indicating a tendency toward overfitting. VGG and ResNet maintain stable validation accuracy, suggesting better generalization. All models achieve high accuracy (>98%) within the first 5 epochs, confirming the effectiveness of feature extraction and classification.
To test real-world performance, we applied data augmentation techniques, including rotation, zoom, flipping, and brightness adjustments. This ensured:

- Increased robustness of models against minor variations in defects.
- Reduction in overfitting by forcing the models to learn invariant features.

CNN's accuracy improved slightly (~0.5%), reducing its misclassification rate for OK parts.

IV. Discussion

In summary, our comparative study of CNN, ResNet18, and VGG16 underscores the potential of deep learning for casting defect detection in submersible pump impellers and similar manufacturing processes. While the baseline CNN achieved strong results, ResNet18 and VGG16 both demonstrated superior performance—often exceeding 98% accuracy—highlighting their capacity for robust feature extraction and classification. These high-accuracy models, when deployed on production lines, can significantly enhance sustainability by reducing the number of defective parts that pass inspection and minimizing the rework or scrapping of otherwise salvageable components. Moreover, incorporating data augmentation proved crucial in promoting model generalization, suggesting that real-world applications will benefit from extensive variability in training samples. Looking ahead, integrating these architectures with IoT-driven sensors and edge computing platforms can further optimize inspection throughput, energy usage, and resource allocation. Ultimately, deep learning-based defect detection not only elevates product quality and operational efficiency but also aligns with broader sustainability objectives in modern casting manufacturing.

References

[1] Swift, K. G., & Booker, J. D. (2013). *Manufacturing process selection handbook.* Butterworth-Heinemann.

[2] Tao, F., Cheng, J., Qi, Q., Zhang, M., Zhang, H., & Sui, F. (2018). Digital twin-driven product design, manufacturing and service with big data. *The International Journal of Advanced Manufacturing Technology, 94*(9), 3563–3576. https://doi.org/10.1007/s00170-017-0233-1

[3] Tabernik, D., Šela, S., Skvarč, J., & Skočaj, D. (2020). Segmentation-based deep-learning approach for surface-defect detection. *Journal of Intelligent Manufacturing, 31*(3), 759–776. https://doi.org/10.1007/s10845-019-01476-x

[4] Sani, A. R., Zolfagharian, A., & Kouzani, A. Z. (2024). Automated defects detection in extrusion 3D printing using YOLO models. *Journal of Intelligent Manufacturing.* Advance online publication. https://doi.org/10.1007/s10845-024-02223-8

[5] Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., ... & Houlsby, N. (2020). *An image is worth 16×16 words: Transformers for image recognition at scale. arXiv preprint* arXiv:2010.11929. https://arxiv.org/abs/2010.11929

[6] Bressanelli, G., Adrodegari, F., Pigosso, D. C. A., & Parida, V. (2022). Towards the smart circular economy paradigm: A definition, conceptualization, and research agenda. *Sustainability, 14*(9), 4960. https://doi.org/10.3390/su14094960

[7] Habibpour, M., Gharoun, H., Tajally, A., Shamsi, A., Asgharnezhad, H., Khosravi, A., & Nahavandi, S. (2021). *An uncertainty-aware deep learning framework for defect detection in casting products. arXiv preprint* arXiv:2107.11643. https://arxiv.org/abs/2107.11643

[8] Ding, Z., Zhao, J., Misra, R. D. K., Guo, F., Xie, Z., Wang, X., ... & Shang, C. (2023). Deep learning-based understanding of defects in continuous casting product. *Metals, 13*(11), 1809. https://doi.org/10.3390/met13111809

[9] Vijay Kumar, V., & Shahin, K. (2025). Artificial intelligence and machine learning for sustainable manufacturing: Current trends and future prospects. *Intelligent and Sustainable Manufacturing, 2*(1), 10002. https://doi.org/10.1016/j.ism.2025.10002

[10] Ripik AI. (2024, November 14). *The future of manufacturing: AI and sustainability working hand in hand. Ripik AI Blog.* https://www.ripik.ai/the-future-of-manufacturing-ai-and-sustainability-working-hand-in-hand/

[11] Ojha, V. K., Goyal, S., & Chand, M. (2025). Developing novel deep learning models for

automated quality inspection in casting. *International Journal of Metalcasting.* Advance online publication. https://doi.org/10.1007/s40962-024-01542-y

[12] Shorten, C., & Khoshgoftaar, T. M. (2019). A survey on image data augmentation for deep learning. *Journal of Big Data, 6*(1), 60. https://doi.org/10.1186/s40537-019-0197-0

[13] Krizhevsky, A., & Hinton, G. E. (2009). *Learning multiple layers of features from tiny images* [Technical report]. University of Toronto. https://www.cs.toronto.edu/~kriz/learning-features-2009-TR.pdf

[14] Kingma, D. P., & Ba, J. (2014). *Adam: A method for stochastic optimization. arXiv preprint* arXiv:1412.6980. https://arxiv.org/abs/1412.6980

[15] He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (pp. 770–778). https://doi.org/10.1109/CVPR.2016.90

[16] Ioffe, S., & Szegedy, C. (2015). Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *Proceedings of the 32nd International Conference on Machine Learning* (Vol. 37, pp. 448–456). PMLR. https://proceedings.mlr.press/v37/ioffe15.html

[17] Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A. C., & Fei-Fei, L. (2015). ImageNet large scale visual recognition challenge. *International Journal of Computer Vision, 115*(3), 211–252. https://doi.org/10.1007/s11263-015-0816-y