Deep vision quality control - an example application in the casting industry

Pedro Rocha[1]
a21200219@isec.pt
Nuno Martins[1,2]
nuno.martins@co.it.pt
Fernando Lopes[1,2]
flopes@isec.pt, flopes@co.it.pt
Luis A. da Silva Cruz[2,3]
lcruz@deec.uc.pt

[1] Polytechnic Institute of Coimbra, Coimbra Institute of Engineering, Coimbra, PT
[2] Instituto de Telecomunicações, Coimbra, PT
[3] Department of Electrical and Computer Engineering, University of Coimbra, Coimbra, PT

Abstract

Ensuring the quality of individual parts during manufacturing is crucial for upholding the integrity of the final products and fostering confidence among customers in the global market.

In this paper, we explore the potential of applying deep learning techniques in a computer vision task, aiming to classify manufactured parts as either defective or non-defective, based on the processing of grayscale images of these parts. We conducted experiments on a publicly available dataset comprising 1300 images of casted submersible pump impellers, utilizing a custom CNN implemented with the Keras framework. We attained a recall rate of 94% for the defective class and a precision of 89% for the non-defective class. Additionally, we discuss qualitative insights gained through the application of Grad-CAM (Gradient-weighted Class Activation Mapping) for better understanding of the model's decision-making process.

1 Introduction

Effective quality control procedures are essential for manufacturing companies, directly influencing their reputation and the reliability of their products. The most common procedures include geometric inspection, which ensures that a product's shape and physical measurements align with specified requirements and tolerances, and surface inspection, which involves comparing attributes like color, texture, brightness, and opacity to predefined standards within acceptable quality ranges. Relying on manual labor for these repetitive tasks introduces variability in quality control outcomes. In our rapidly evolving industrial landscape, time is an increasingly valuable resource. To meet the demands of 24/7 manufacturing, technology-driven solutions have become indispensable, benefiting both employers and employees.

The casting industry is a pivotal segment of manufacturing, involving the production of metal components through the casting process. In this process, liquid material, typically metal, is poured into a mold with a hollow cavity of the desired shape. The material cools and solidifies, resulting in a part or product with the intended shape [1]. During this process, irregularities or imperfections can occur, significantly impacting product quality and reliability. Common casting defects, include inclusions, porosity, runout, flashing, cold shuts, blowholes, pinholes, burrs, shrinkage defects, problems with mold materials, and issues with metal pouring. Quality control in the casting industry is of utmost importance, ensuring that the final products meet required standards and are defect-free. Traditionally, quality inspection has relied on manual inspections by human inspectors. However, this manual process is time-consuming and susceptible to inaccuracies due to human limitations.

The limitations of manual inspection can lead to the rejection of entire orders due to defects in just a few products, resulting in significant financial losses for manufacturing companies. To address these challenges, the industry is increasingly turning to automation and technology [2]. One promising approach is the use of deep learning-based classification models to automate the inspection process. These models can analyze images of cast products and identify defects with a high degree of accuracy, reducing the reliance on manual inspection and enhancing overall quality control. By implementing deep learning classification models, manufacturing companies can improve the efficiency and accuracy of their quality control processes, reducing the likelihood of defective products reaching the market. Ultimately, this approach can enhance their bottom line by minimizing losses due to rejected orders.

2 Literature review

In the study made by Bart De Ketelaere *et al.* [3], it is emphasized that with the the arise of computational power to leverage deep learning algorithms and multispectral or hyperspectral imaging systems, there is the potential to extract richer chemical information from object samples beyond the visible light spectrum. This capability enables the visualization of aspects that are otherwise imperceptible to the human eye, such as food packages sealing conditions.

Matthias Becker [4] explored the concept of controlling the mold heating and cooling operation in the light metal die casting process through real-time monitoring, using infrared cameras and low-end embedded process control computers runing a Convolutional Neural Network (CNN).

Using cast product images for quality inspection Seokju Oh *et al.* [5] achieved an increase of 7.6% (from 90.92% to 98.58%) in the F1-score performance by employing data augmentation via a Convolutional Autoencoder (CAE) in their proposed CNN. After that, Seokju Oh, along with Juyong Park *et al.* [6] evaluated the classification performance on the Xception deep leaning model, improving defect detection by proposing the use a Wavelet Transform Denoise algorithm after applying a Gaussian noise filter to the input.

Max Ferguson *et al.* [7] developed XnetV2, a defect classifier based on the Xnet architecture, to localize casting defects in X-Ray images. They conducted a comprehensive study exploring the use of several state-of-the-art object detectors, including Faster R-CNN, R-FCN, and SSD. Decoupling the feature extraction layer from the object detection architecture allowed them to assess the performance of each object detection method with different feature extractors. Using an adapted version of the Faster R-CNN architecture, they achieve a mAP of 0.921 on test.

3 Methodology

In a production line, components on conveyors undergo inspections via machines equipped with detection photocells, that send a trigger signal to the camera sensor [8]. Although minor positional variations in objects may occur, overall light conditions remain stable. After image acquisition comes the subsequent stage of image processing. Following that, the final phase involves a verdict decision, which often translates into product rejection, accomplished by changing the item's trajectory towards a dedicated conveyor circuit for potential reintegration. This particular application is characterized by its demand for speed, given its real-time inspection capabilities. While precision is paramount to prevent wastage of production resources, achieving high recall is critical to guarantee product quality and integrity throughout the manufacturing process. In this work, we evaluate the applicability of deep vision in this context by implementing a Convolutional Neural Network (CNN) to detect casting defects in images of submersible pump impellers, as exemplified in Figure 1.

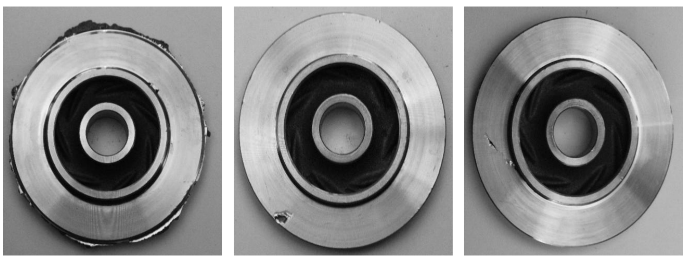

Figure 1: Examples of Defective Pump Impellers

3.1 Dataset

We used a publicly available dataset from Kaggle, provided by Ravirajsinh Dabhi with Pilot Technocast colaboration [9]. This dataset comprises 7348 grayscale images with dimensions of 300x300 pixels, which have already undergone data augmentation. Additionally, there is an original dataset consisting of 1300 grayscale images with a resolution of 512x512 pixels and no augmentation. Although the dataset comprises various types of defects, the critical aspect for quality control lies in identifying the absence of them, making this problem a binary classification with just two classes: 'Defective' and 'OK.'

Being required to split the dataset in three, for training, validation, and testing, each split containing unique images, it is important to perform data augmentation only after splitting. Furthermore, data augmentation is typically applied exclusively to the training dataset. After the first experiments to be described in the next section, where we used the augmented 300x300 dataset, we opted to use the original 512x512 dataset and split it into proportions of 80%, 10%, and 10% for training, validation, and testing, respectively. Furthermore, we applied data augmentation on the original dataset before training by incorporating random operations such as horizontal and vertical flips, rotations, zooming (limited to 10%), and adjustments in brightness and contrast (limited to 15%).

3.2 Architecture

We began our experiments with a custom model comprising five convolutional layers and five dense layers with the 300x300 dataset. In fewer than ten epochs, it achieved an impressive 98% accuracy on all three dataset splits, similar to previous works [5, 6]. However, this raised doubts about dataset integrity, reinforced by the fact that augmentation and split operations were not clear in the dataset characterisation. GradCam revealed that the model was not assigning significant decision weight to pixels with actual defects. Instead, it seemed to encode each image randomly, achieving high scores due to dataset integrity issues. Subsequently, we decided to experiment with the 512x512 dataset, yielding poor results. We then attempted to integrate a pretrained ResNet50 model, fine-tuned only in the last dense layers. However, we encountered challenges during training, with the model failing to converge. Ultimately, we decided to proceed with a custom model featuring a normalization layer scaling to the range [-1, 1], six Conv2D layers, where the first half had 32 filters and the second half had 64. The first layer utilized a 5x5 kernel and a stride of 3, while the subsequent layers used a 3x3 kernel and a stride of 2. After flattening, we included three dense layers with sizes of 64, 128, and 256, respectively, followed by the final sigmoid dense layer. The motivation for widening the layers in the dense block was to reduce the complexity.

4 Results and evaluation

In this section, we evaluate our custom model with the original 512x512 dataset. By setting the classification threshold to 0.50, we achieved the values for the metrics detailed in Table 1. The confusion matrix presented in Figure 2 (left) shows the distribution of the predictions (TN, FN, FN, FP) in the split dataset for testing. To provide insights into the trade-off between correctly identifying true positives and incorrectly classifying false positives across various classification thresholds, we present the Receiver Operating Characteristic (ROC) curve in Figure 2 (right), where our model achieved an Area Under the ROC Curve (AUC-ROC) of 97%.

Table 1: Test inference performance metrics (threshold=0.50)

Class	Precision	Recall	F1-Score	# samples
OK	0.89	0.83	0.86	48
Defect	0.91	0.94	0.92	85
Avg / Total	0.90	0.90	0.90	133

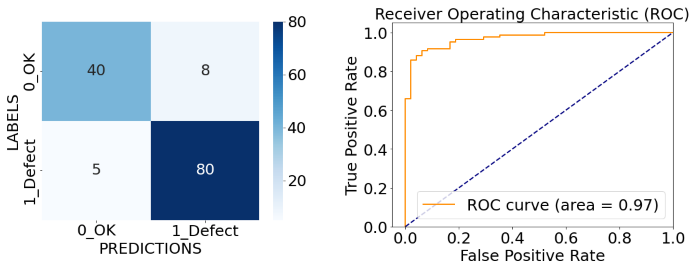

Figure 2: Confusion matrix and ROC curve

Grad-CAM (Gradient-weighted Class Activation Mapping) is a technique used to visualize the regions within an image that most significantly contributed to the final classification decision. This is achieved by computing the gradients of the decision scores at each pixel in the image with respect to the final classification outcome. These gradients are then weighted and aggregated across a specific convolutional layer's depth. In Figure 3, we showcase the results obtained using the Grad-CAM across all six Conv2D layers on two images that were misclassified by our model.

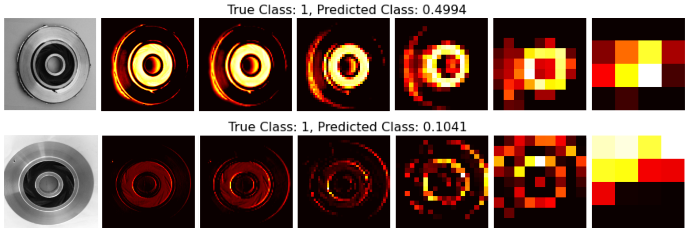

Figure 3: GradCam results of two missclassified examples

In these images, as well as in others, we can observe that the actual defect is indeed highlighted. However, it is essential to note that the defect may not always be the brightest element within the image. Instead, our observation suggests that the model's classification decision is likely influenced by the features of the central ring, which appears to play a significant role in the final decision.

5 Conclusions and future work

In this study, we assessed the effectiveness of a custom CNN deep learning model in classifying parts from images acquired in a real-world automated inspection line, deciding on their acceptability in terms of quality. The obtained results are good, but still with potential for improvement. In particular, the feature extraction could be enhanced by increasing the quality of the datasets. As part of our future work, we plan to collaborate with a local industry to gain access to a substantial amount of labeled data. In addition, we plan to design an unsupervised learning pre-processing layer to act as a binary mask, bounding the ROI, and making the model decision focus on the actual features rather than noise.

References

[1] Rajesh Rajkolhe, J. Khan, Shri Sant, and Gajanan Maharaj. Defects, Causes and Their Remedies in Casting Process: A Review. 2014.

[2] Natalia O. Kurdyukova, Iurii V. Ognev, and Kymbat B. Kunanbayeva. Implementation of Machine Vision and Other Automated Systems to Ensure the Collection of Representative Statistics of Process Control and Quality Management at Industrial Enterprises. In *2021 IT&QM&IS*, pages 309–312, September 2021.

[3] Bart De Ketelaere, N. Wouters, I. Kalfas, R. Van Belleghem, and W. Saeys. A fresh look at computer vision for industrial quality control. *Quality Engineering*, 34(1):152–158, January 2022.

[4] Matthias Becker. Quality control of a light metal die casting process using artificial neural networks. In *2009 IEEE Int. Conference on Computational Cybernetics (ICCC)*, pages 113–117, January 2009.

[5] Seokju Oh, Jaegyeong Cha, Donghyun Kim, and Jongpil Jeong. Quality Inspection of Casting Product Using CAE and CNN. In *2020 4th International Conference on Imaging, Signal Processing and Communications (ICISPC)*, pages 34–38, October 2020.

[6] Juyong Park, Seokju Oh, and Jongpil Jeong. A Defect Detection Model for Casting Product using Wavelet Transform Denoise and Xception. In *2022 24th International Conference on Advanced Communication Technology (ICACT)*, pages 247–251, Feb. 2022.

[7] Max Ferguson, Ronay Ak, Yung-Tsun Tina Lee, and Kincho H. Law. Automatic localization of casting defects with convolutional neural networks. In *2017 IEEE International Conference on Big Data (Big Data)*, pages 1726–1735, December 2017.

[8] Longfei Zhou, Lin Zhang, and Nicholas Konz. Computer Vision Techniques in Manufacturing. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 53(1):105–117, January 2023.

[9] Ravirajsinh Dabhi. Casting product image data for quality inspection, 2020. https://www.kaggle.com/datasets/ravirajsinh45/real-life-industrial-dataset-of-casting-product.