Deep Learning for Robust Manufacturing: From Quality Control to Predictive Maintenance

Paula Mielgo[1], Anibal Bregon[2] and Miguel A. Martínez-Prieto[3]

[1,2,3] *Department of Computer Science, University of Valladolid, Valladolid, Spain*
paula.mielgo@uva.es, anibal.bregon@uva.es, miguelamp@uva.es

Abstract

In today's global manufacturing landscape, companies are required to balance speed, precision, and sustainability, thereby making intelligent, data-driven solutions a necessity. The convergence of Industry 4.0, cyber-physical systems and artificial intelligence technologies is leading to a new paradigm known as smart manufacturing, where the effective use of collected data can increase productivity, efficiency, and quality. This research explores the potential of deep learning to enhance industrial productivity by leveraging automated quality control of manufactured components and predictive maintenance of systems. Thus, this thesis focuses on two main objectives within an industrial context: (OB1) the development of models to enhance the quality control process, and (OB2) the development of models to implement a predictive maintenance strategy. These objectives are approached through three expected contributions. First, a quality control model for thermal images aligned with factory requirements (C1). Second, a contrastive learning model for anomaly detection in multiview images (C2). Third, a predictive maintenance model for die-casting molds (C3). Initial results are starting to show the advantage of these contributions to improve productivity in smart factories.

1. Introduction

With the growing complexity of global manufacturing, companies face increasing pressure to improve efficiency, product quality, and sustainability while reducing costs. Industry 4.0 aims to digitally transform manufacturing and production systems through the integration of advanced technologies, including cyber-physical systems and artificial intelligence. This transformation enables the development of smart factories where data is leveraged at every stage of the manufacturing process to optimize decision-making and operational performance, thus enhancing productivity.

Paula Mielgo et al. This is an open-access article distributed under the terms of the Creative Commons Attribution 3.0 United States License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

In recent years, the field of artificial intelligence has experienced a revolution with the potential to significantly transform production processes. Recent developments have been successfully applied to production optimization, quality control, and predictive maintenance. The implementation of Machine Learning (ML) algorithms within industrial production processes has been demonstrated to enhance quality and reduce production time. Several literature reviews highlight the application of ML techniques in manufacturing operations, such as planning procedures (Usuga Cadavid, Lamouri, Grabot, Pellerin, & Fortin, 2020) and quality control (Peres, Barata, Leitao, & Garcia, 2019) (Patel & Jokhakar, 2016). In addition, the advent of deep learning (DL) has enabled significant progress in computer vision, enhancing a variety of application domains, including manufacturing. Architectures such as Convolutional Neural Networks (CNNs) and Vision Transformers (ViTs) have proven effective in industrial image and video processing for quality control (EL Ghadoui, Mouchtachi, & Majdoul, 2023) (Villalba-Diez et al., 2019) (Cumbajin et al., 2023) and predictive maintenance (Hassan, Steinnes, Gustafsson, Løken, & Hameed, 2023). Through the analysis of visual data, these models can detect defects, monitor processes, and anticipate failures, tasks that are traditionally challenging for human inspectors or rule-based systems.

While significant progress has been reported in the literature, several challenges remain. Quality control models often have problems dealing with highly imbalanced datasets and face difficulties in real-time processing. Moreover, predictive maintenance models in industrial environments must be robust and capable of forecasting degradation trends from limited and noisy historical data.

This thesis addresses these critical challenges by enhancing product and process reliability in smart manufacturing environments through the unification of quality control and predictive maintenance strategies under a common framework. This integrated approach enables a complete solution for maintaining quality, reliability, and efficiency in manufacturing systems. Both components rely on deep learning to support intelligent, data-driven decision-making that safeguards both product quality and process integrity.

2. PROBLEM STATEMENT

Ensuring the quality of manufactured products is a critical challenge in modern industrial processes. The early identification of defective parts on the production line results in improved factory productivity and reduced manufacturing costs. However, the inspection of product quality using manual methods or traditional rule-based systems can be time-consuming, limited in detect subtle anomalies, and prone to human error, thereby reducing process efficiency. The application of deep learning techniques has proven effective adresing these limitations, particularly in anomaly detection. However, two challenges hinder the effective deployment of deep learning models in industrial environments. First, data imbalance represents a significant obstacle, since datasets often contain a small number of defective instances compared to non-defective ones. This imbalance frequently leads to biased models that perform poorly on minority classes, thereby limiting the overall effectiveness of the system. Second, real-time constraints must be considered. Models must be capable of processing new data in real time to support the decision-making process. Otherwise, the quality assurance test could slow down the manufacturing process, resulting in a reduced production rate. An additional challenge arises in predictive maintenance. Machine components experience progressive wear and degradation over time, resulting in increased defect rates and potential downtime. Current monitoring systems struggle to predict this degradation accurately, thereby limiting the ability to schedule maintenance interventions.

The purpose of this thesis is to address these gaps by developing deep learning models that:

- Enhance defect detection performance in quality control tasks, even with highly imbalanced datasets.
- Enable real-time decision-making for industrial operations.
- Forecast assets degradation trends to support effective predictive maintenance strategies.

In order to to visualize and align the core elements of a dissertation, we have used the CQOCE diagram (Prieto, 2019). The CQOCE (Context, Question, Objectives, Contributions, and Evaluation) diagram is a powerful synthesis and communication tool for doctoral research. Accordingly, Figure 1 presents the CQOCE diagram corresponding to the topic of this thesis, complementing the information discussed in the following sections.

3. RESEARCH QUESTIONS

To address the challenges presented in the problem statement, the following main research question is formulated.

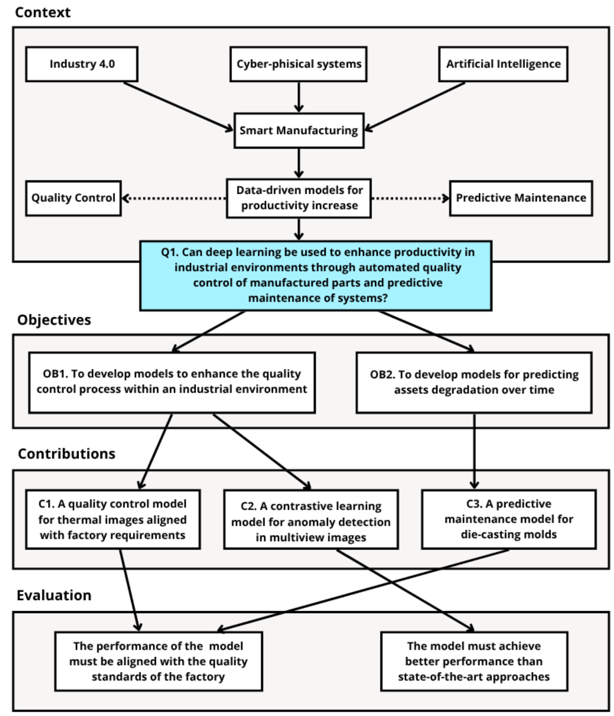

Figure 1. CQOCE diagram.

Q1. Can deep learning be used to enhance productivity in industrial environments through automated quality control of manufactured parts and predictive maintenance of systems?

This main question can be addressed through the following sub-questions.

Q2. Can deep learning models be designed to effectively detect subtle product defects in industrial manufacturing processes, particularly in the presence of imbalanced data?

Q3. Can deep learning techniques be applied to accurately predict the progressive degradation of critical production assets to enable predictive maintenance strategies?

In order to respond to these questions, two objectives have been formulated.

OB1. To develop models to enhance the quality control process by effectively identifying defective parts.

OB2. To develop models for predicting assets degradation over time, contributing to the implementation of effective predictive maintenance strategies.

4. EXPECTED CONTRIBUTIONS

The thesis is focused on an industrial context. To this end, two datasets were selected. On the one hand, a proprietary dataset composed of thermal images captured after a die-casting process (Figure 2a). It consists of pairs of RGB images of the mold used for each manufactured part, rather than of the part itself. Following the casting stage, every part undergoes three quality assurance procedures: an internal quality control system of the die-casting machine, a visual inspection conducted by an operator, and a leak test. All images in the dataset correspond to parts that were classified as non-defective by the first two tests. Consequently, the final label assigned to each image is based only on the outcome of the leak test: parts that pass it are labeled as OK, while those that fail it are labeled as NK. It is hypothesized that defective parts, having passed the initial controls, present only minor porosity defects, potentially imperceptible to the human eye. Initial experiments using CNN and ViT architectures exhibited poor results (Mielgo et al., 2024). In particular, with this solution, the high rate of defective parts misclassified as non-defective that would passed downstream from the assembly line would make the system unsuitable for industrial deployment. Therefore, we want to develop a novel architecture capable not only of detecting defective parts with higher accuracy but also of aligning performance metrics with the strict requirements of factory operations. Furthermore, as part of our prognostic objectives, we also want to estimate mold degradation over time throughout the manufacturing process. This would contribute to extending the service life of the mold components. On the other hand, a large-scale public dataset composed of real-world multi-view images for industrial anomaly detection, the Real-IAD dataset (Wang et al., 2024) (Figure 2b). Real-IAD provides 150K high-resolution images across 30 object categories captured from multiple angles under real conditions, allowing robust benchmarking for anomaly detection tasks. It tries to better reflect real-world inspection scenarios including also a wider range of defect sizes and types, making it more challenging than existing benchmarks. A multiview architecture based on a contrastive loss function (Chen, Kornblith, Norouzi, & Hinton, 2020) is expected to be developed to perform supervised classification on this dataset. To the best of our knowledge, this would represent the first application of contrastive learning to an industrial multiview dataset for anomaly detection.

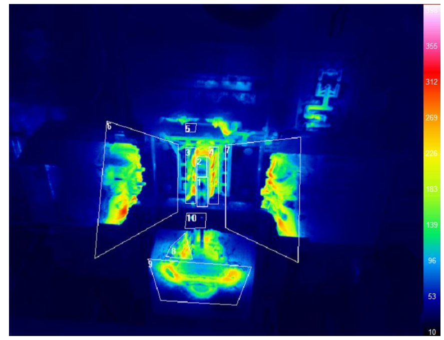

(a) Thermal dataset

(b) Real-IAD dataset

Figure 2. Selected datasets

5. PROPOSED RESEARCH PLAN

The research project was initiated in the fall of 2023 and is planned to be concluded in the fall of 2027. The work is organized according to an iterative and incremental methodology. On the one hand, preliminary results corresponding to each research question will be subjected to validation by the scientific community through symposiums or conference publications. Through these channels, we expect to receive valuable feedback that will guide the direction of our research. In addition, this allows us to verify the validity of our proposal in terms of interest and quality before conducting a more extensive research study. On the other hand, the final results will be oriented towards journal publications and international conferences, as the relevance of the approach has already been validated and is expected to provide significant results. Finally, all the knowledge accumulated during the thesis period will be consolidated in the thesis, which will be developed throughout the entire PhD program.

The main activities planned for each contribution in the thesis are listed below. Figure 3 illustrates the timeline over the 16 planned quarters.

C1. A quality control model for thermal images aligned with factory requirements.

A1. Review of the main state-of-the-art techniques employed for anomaly detection in industrial contexts.

A2. Data preparation and cleaning of the thermal image dataset.

A3. Implementation and training of a quality control model for anomaly detection in thermal images.

A4. Model evaluation.

C2. A contrastive learning model for anomaly detection in multiview images.

A5. Review of the main state-of-the-art contrastive learning architectures.

A6. Implementation and training of a multiview contrastive learning architecture for anomaly detection.

A7. Model evaluation.

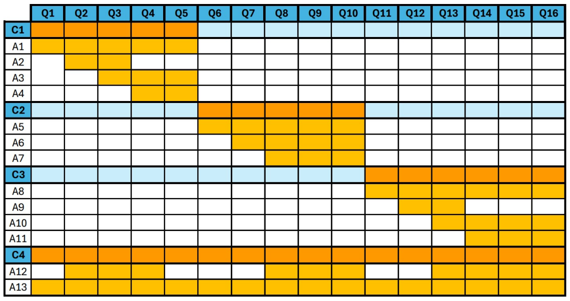

Figure 3. Research plan timeline.

C3. A deep learning model for predicting mold degradation over time, contributing to the implementation of effective predictive maintenance strategies in die-casting processes.

A8. Review of state-of-the-art predictive maintenance techniques applied in industrial environments.

A9. Data preparation and cleaning of the thermal image dataset.

A10. Implementation and training of a prognostic model for mold degradation detection.

A11. Model evaluation.

C4. Scientific dissemination.

A12. Manuscript preparation and submission to international conferences and journals.

A13. Thesis writing and defense preparation.

6. CURRENT PROGRESS AND PRELIMINARY RESULTS

The initial contribution block C1 has been completed, with preliminary results published at the 2024 PHM conference (Mielgo et al., 2024). The research was centered on the development of a quality control model for thermal images in a die-casting machine. The proposed architecture employs two models and two thresholds to perform a partial classification, remaining some instances unlabeled. Consequently, we prioritize the number of parts that can avoid extra quality tests while ensuring the confidence of the prediction. This approach outperforms the baseline, achieving a 22% reduction in the number of manufactured parts that would require additional quality tests. Additionally, a more advanced architecture is under review for a journal publication.

Work is ongoing on the second contribution block (C2), with preliminary results to be submitted to the 2025 PHM Conference. A multiview contrastive learning approach shows promising improvements over traditional single-view methods.

ACKNOWLEDGMENT

Paula Mielgo's work has been funded by UVa 2023 predoctoral contracts, co-financed by Banco Santander. This work has been funded by the Spanish Ministerio de Ciencia e Innovación under grant PID2021-126659OB-I00. We acknowledge HORSE Powertrain for granting us access to the thermographic images dataset, and Javier Moral and Daniel Lopez for the help and support.

REFERENCES

Chen, T., Kornblith, S., Norouzi, M., & Hinton, G. (2020). A simple framework for contrastive learning of visual representations. In *International Conference on Machine Learning* (pp. 1597–1607).

Cumbajin, E., Rodrigues, N., Costa, P., Miragaia, R., Frazão, L., Costa, N., … Pereira, A. (2023). A real-time automated defect detection system for ceramic pieces manufacturing process based on computer vision with deep learning. *Sensors*, *24*(1), 232.

EL Ghadoui, M., Mouchtachi, A., & Majdoul, R. (2023). Intelligent surface roughness measurement using deep learning and computer vision: a promising approach for manufacturing quality control. *The International Journal of Advanced Manufacturing Technology*, *129*(7), 3261–3268.

Hassan, M. U., Steinnes, O.-M. H., Gustafsson, E. G., Løken, S., & Hameed, I. A. (2023). Predictive maintenance of norwegian road network using deep learning models. *Sensors*, *23*(6), 2935.

Mielgo, P., Bregon, A., Alonso, C. J., López, D., Martínez-Prieto, M. A., & Pulido, B. (2024). A Deep Learning Solution for Quality Control in a Die Casting Process. In *Annual Conference of the PHM Society* (Vol. 16).

Patel, S., & Jokhakar, V. N. (2016). A random forest based machine learning approach for mild steel defect diagnosis. In *2016 IEEE Int. Conf. on computational intelligence and computing research (ICCIC)* (p. 1-8).

Peres, R. S., Barata, J., Leitao, P., & Garcia, G. (2019). Multistage quality control using machine learning in the automotive industry. *IEEE Access*, *7*, 79908-79916.

Prieto, L. P. (2019). *A Happy PhD — PhD tool: The CQOCE diagram.* https://ahappyphd.org/posts/cqoce-diagram/. (Accessed: 2025-05-07)

Usuga Cadavid, J. P., Lamouri, S., Grabot, B., Pellerin, R., & Fortin, A. (2020). Machine learning applied in production planning and control: a state-of-the-art in the era of industry 4.0. *Journal of Intelligent Manufacturing*, *31*, 1531-1558.

Villalba-Diez, J., Schmidt, D., Gevers, R., Ordieres-Meré, J., Buchwitz, M., & Wellbrock, W. (2019). Deep learning for industrial computer vision quality control in the printing industry 4.0. *Sensors*, *19*(18), 3987.

Wang, C., Zhu, W., Gao, B.-B., Gan, Z., Zhang, J., Gu, Z., … Ma, L. (2024). Real-iad: A real-world multi-view dataset for benchmarking versatile industrial anomaly detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition.*