euspen's 26th International Conference & Exhibition, Krakow, PL, June 2026
www.euspen.eu

Deep learning-driven visual inspection methodology for automated defect detection and dimensional assessment in die-cast parts

Sofia Catalucci and Enrico Savio

University of Padua, Department of Industrial Engineering (DII), Italy

sofia.catalucci@unipd.it

Abstract
Defect detection is critical across industrial manufacturing processes such as casting, moulding, machining, and additive manufacturing, as imperfections (e.g., burrs, porosities, surface discontinuities) can compromise assembly, performance, and structural reliability. When defects exhibit a three-dimensional morphology, conventional image-based inspection becomes less effective and human visual inspection struggles to deliver consistent, objective decisions. This work presents a metrologically validated pipeline that employs optical acquisition of 3D point clouds from a machine vision system by means of laser triangulation, compares them with a reference point cloud acquired from a conforming part, and converts the resulting deviation maps into images for automated deep learning-driven defect localisation, followed by skeletonization-based dimensional assessment. Validation is performed against reference measurements acquired with a multisensor coordinate measuring machine. Automotive die-cast aluminium components are used as a case study to demonstrate robust burr detection and dimensional quantification under industrial constraints, supporting integration into in-line quality control workflows.

Defect detection; deep learning; dimensional measurement; metrological validation; die-casting

1. Introduction

In modern manufacturing, surface defect detection is pivotal to ensuring the functionality and reliability of industrial components, particularly in high-volume production, where even small surface imperfections can lead to assembly issues, failures, or downstream rework [1]. In this context, three-dimensional (3D) surface defects (i.e., geometrical imperfections such as protrusions, discontinuities, or burrs) are particularly critical because their detectability depends strongly on scale, morphology, and local surface conditions. Traditionally, detection of surface defects relies on human visual inspection (HVI) supported by simple tools. While HVI remains widespread, it is inherently affected by operator subjectivity and fatigue, which limit repeatability and robustness. To mitigate these issues, automated machine vision systems (MVS) are increasingly deployed in-line; however, their effectiveness still depends critically on illumination setup and on the complexity of part geometry and surface characteristics. Changes in lighting, surface reflectivity, and viewing angle can significantly alter defect visibility and detection reliability.

Recently, deep learning (DL) approaches have improved defect detection performance; however, many solutions provide qualitative outputs and rarely include metrological validation. This work introduces a DL-based inspection pipeline that couples 3D surface defect localisation with dimensional assessment via skeletonization [2]. The approach is demonstrated on burr detection and quantification on automotive die-cast aluminium components under industrially representative conditions.

2. Materials and Methods

The proposed pipeline adopts supervised transfer learning, using a convolutional detector initialised from pre-trained weights and fine-tuned on pre-labelled data to localise surface defects. The detector operates on two-dimensional (2D) deviation maps obtained by comparing 3D point clouds acquired via laser triangulation on die-cast parts against a high-density scan of a conforming part. After registration in a common reference frame, the resulting deviation maps are converted into images, augmented to mimic industrial variability, and used as input to the network. For each detected region, defect dimensions are estimated through a skeletonization-based step that extracts the defect medial axis and related geometric descriptors. The methodology is validated against reference measurements acquired using a multisensor coordinate measuring machine (CMM). The workflow is shown in Figure 1.

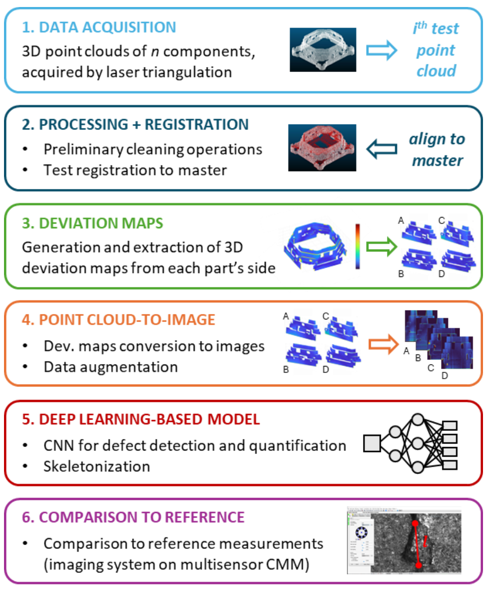

Figure 1. Workflow for defect detection and quantification.

2.1. Case study and data acquisition

The present work focuses on detecting burrs (i.e., raised sharp edges) on aluminium alloy AlSi12(Fe) housings for automotive long-range radar sensors [2]. On these components, residual burr size typically ranges from 0.4 to 2 mm, originating from both the moulding stage and the subsequent shearing operations used to remove the sprue and casting attachments. Parts were measured using a laser triangulation system Nikon MCAx, generating point clouds with an average point sampling of 100 µm. This sampling supports reliable visualisation of defects ≥ 1 mm and provides sufficient detail for imperfections of approx. 0.6 mm. Each scan was registered to a reference point cloud acquired from a conforming part with higher density (10 µm average point sampling). Registration was performed by means of best-fit algorithm in the software CloudCompare. Table 1 reports both instruments' specifications.

Table 1. Instruments specifications

	Technology	Sampling	Accuracy*	Point count
Reference optical setup	Laser-based	10 µm	±0.2 µm	approx. 1,100,000
AACMM Nikon MCAx	Laser-based	100 µm	±10 µm	approx. 800,000

*stated by the manufacturer (term not used in the strict VIM sense)

2.2. Image preparation and deep learning model

After alignment, 3D deviation maps representing local surface height differences with respect to the master geometry were obtained. The four lateral faces of each component were segmented (Sides A–D), and the resulting deviation maps were projected onto reference planes and exported as 2D colour-encoded images. These images were independently normalised for input to the detector, while missing-data regions were handled via masking. From 92 parts, 368 images were generated and a balanced subset (85 defective, 75 non-defective) was selected to reduce class bias. Data preparation and augmentation were performed in Roboflow (online tool for computer vision tasks) to simulate realistic in-line inspection variability (exposure ±10%, brightness and saturation ±15%, rotation ±5°), producing a working dataset of 400 images. A bounding box annotation operation was performed on a subset of defect cases. All images were resized to 640 × 640 pixels (without padding), and the dataset was split into train/validation/test (70/20/10), keeping validation/test restricted to non-augmented images to ensure unbiased evaluation. The network adopted in this work was YOLOv8 (Ultralytics, Python API).

2.3. Defect detection, dimensional inspection and validation

After training, the DL model in its best configuration (based on the metrics defined in [3]) was applied to the unlabelled images. For each image predicted as defective, defect localisation was obtained from a) the predicted bounding box, providing position and an initial estimate of the defect extent; b) dimensional assessment was then refined by applying a skeletonization algorithm to the detected region. The total length L_{sk} of each detected defect was computed as the sum of the lengths l_{sk_i} of the n skeleton segments that compose it, as in equation (1).

$$L_{sk} = \sum_{i=1}^{n} l_{sk\_i} \qquad (1)$$

All pixel-based measurements were converted to metric units using a conversion factor equal to 0.046 mm/pixel, derived from the part height (29.4 mm) and the image resolution (640 pixels).

Reference measurements were acquired using a multisensor CMM ZEISS O-Inspect equipped with an imaging probing system (MPE = 1.9 + L/150 µm for 2D length measurements), using a repeatable clamping fixture and two alignments to inspect all four lateral faces of the part. A dimensional tolerance of 0.6 mm was set as the limit beyond which any detected protrusion is deemed non-conforming, and the part should be rejected. For validation, burr size was quantified by averaging five measurements of the distance $\bar{l}$ between two points selected by the user at the burr endpoints. Examples are shown in Figure 2.

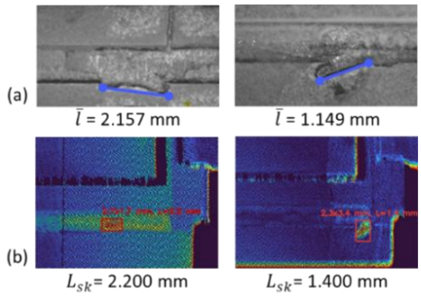

Figure 2. (a) Imaging lengths, and (b) DL-driven skeletonization.

3. Results and discussion

The YOLOv8 detector was trained across five configurations (5, 10, 20, 50, and 100 epochs). Model selection was based on the best validation epoch, rather than the final epoch, to capture peak generalisation and limit late-stage overfitting. Performance improved from 5 to 50 epochs, with the 50-epoch configuration achieving the best overall results, as 100 epochs indicated increased selectivity and potential overfitting. Bounding box initial estimate of the defect extent showed an average deviation of 0.30 mm from the imaging sensor reference for burrs of 1.5-2 mm and above, but it systematically overestimated defects below 1.5 mm. For this reason, estimates were complemented with skeletonization. Compared to the imaging reference, the skeleton-based length showed a mean deviation of 0.29 mm. Missed detections were attributed to insufficient point density, which affected the deviation maps. Skeletonization tended to slightly overestimate sub-millimetre defects, while being more conservative for larger ones.

4. Conclusions and outlook

This work proposes a metrologically validated pipeline for automated burr detection and dimensional assessment on die-cast parts. YOLOv8 was applied to deviation maps converted into images derived from registered 3D scans, and defect length was estimated via bounding box and skeletonization. Dimensional results were validated against reference measurements. From an industrial perspective, the method can support HVI within quality control systems. Current limitations mainly relate to point cloud quality and reduced robustness for sub-millimetre defects. Future work will focus on extending to additional defect types/processes, and exploring in-line data collection and active learning to improve generalisation and scalability.

Acknowledgements

The project (23IND08 DI-Vision) has received funding from the European Partnership on Metrology, co-financed from the European Union's Horizon Europe Research and Innovation Programme and by the Participating States.

References

[1] Catalucci S, Thompson A, Piano S, Branson DT and Leach R 2022 Optical metrology for digital manufacturing: a review *Int. J. Adv. Manuf. Technol.* **120** 4271-4290

[2] Catalucci S and Savio E 2026 Metrological validation of a deep learning pipeline for in-line detection and dimensional quantification of three-dimensional surface defects *CIRP J. Manuf. Sci. Technol.* **64** 107-119

[3] Mullany B, Savio E, Haitjema H and Leach R 2022 The implication and evaluation of geometrical imperfections on manufactured surfaces *CIRP Ann.* **71** 717-739