Additive Manufacturing Letters 17 (2026) 100360

Contents lists available at ScienceDirect

Additive Manufacturing Letters

journal homepage: www.elsevier.com/locate/addlet

Short Communication

Multi-camera fault detection in fused filament fabrication printing

Shanthalakshmi Kilambi [a,b,1], Aster Tournoy [a,c,1], Muhamad Amani [a], Jovana Jovanova [c], Baris Caglar [b], Kunal Masania [a,*]

[a] Shaping Matter Lab, Faculty of Aerospace Engineering, Delft University of Technology, Kluyverweg 1, 2629 HS Delft, Netherlands
[b] CAMP Lab, Faculty of Aerospace Engineering, Delft University of Technology, Kluyverweg 1, 2629 HS Delft, Netherlands
[c] JAM Lab, Faculty of Mechanical Engineering, Delft University of Technology, Mekelweg 2, 2628CD Delft, the Netherlands

ARTICLE INFO

Keywords:
Fused filament fabrication
Fault detection
Convolutional neural networks
Data fusion

ABSTRACT

Fused filament fabrication is a popular extrusion 3D printing technology because of its affordability and accessibility. However, the approach often suffers from printing errors that result in wasted time, materials and energy. Convolutional neural networks can be trained to recognise a wide spectrum of printing anomalies from image data in real time, but past work has been limited to a few defect classifications at a time. Here, we introduce a fault detection system, designed to identify a range of errors without interrupting the printing process. Real-time detection is achieved using a pre-trained image recognition and pattern recognition convolutional neural network (CNN) with two mounted cameras on the print bed and a nozzle camera. Two CNN models are developed to classify images into common 3D printing errors for the two camera systems. The nozzle camera model achieves a high validation accuracy of 97.7%. The side camera model achieves comparable performance with a validation accuracy of 97.6%. To integrate the two CNNs into one unified system, a logic-based priority framework was used to improve reliability beyond individual model accuracies by resolving conflicting predictions and leveraging complementary viewing angles from both camera types to detect a broader range of defects. The data fusion framework identifies 12 common errors and has significantly improved the robustness of error classification, in-situ and in real-time, with inference times as small as 220 milliseconds. The results demonstrate the feasibility of a robust multi-input fault detection system to advance the reliability of extrusion 3D printing.

Introduction

Additive manufacturing (AM) or 3D printing transforms computer-aided design (CAD) files into physical objects by systematically joining material layers[1,2]. This method enables the fabrication of custom and complex structures with significantly less material waste than conventional manufacturing, which typically uses subtractive methods[3]. Among AM techniques, Fused Filament Fabrication (FFF) is the most common, valued for its ease of use and material compatibility[3–5]. The process of FFF involves melting a thermoplastic filament and extruding layer by layer, where the layers solidify to build a 3D structure[3,6–10]. While FFF is cost-effective, it suffers from low resolution[11], poor surface quality and a high frequency of defects[4,12].

The increasing complexity of additively manufactured materials underscores the critical need for this real-time defect detection to enable the efforts aiming to address the defects in situ. Advanced applications rely on fabricating parts where material and structural properties are programmed directly during printing. This makes process control essential, as the manufacturing process itself dictates the material's final properties and performance[13]. For instance, the mechanical properties of bio-inspired composites from liquid crystal polymers are determined by the molecular alignment during extrusion[14]. Any deviation can cause severe degradation of the component's strength. Similarly, in multi-material printing of smart, actuated structures, a small defect like poor interfacial bonding can render an entire device non-functional[15,16]. In the development of engineered living materials, printing errors can lead to inconsistencies in growth, resulting in uncertain material properties or functional behaviour[17]. For these applications, post-process inspection is insufficient, making in-situ monitoring an indispensable tool for ensuring both structural and functional integrity.

While these advanced materials represent the frontier of AM, even conventional FFF printing is plagued by defects[1,18–20]. These issues

* Corresponding author.
E-mail address: k.masania@tudelft.nl (K. Masania).
[1] Authors contributed equally

https://doi.org/10.1016/j.addlet.2026.100360
Received 17 October 2025; Received in revised form 25 January 2026; Accepted 25 January 2026
Available online 26 January 2026

are especially prevalent among inexperienced users, who face failure rates of approximately 20%[21]. This high failure rate stems from the open-loop nature of most 3D printers. To address this, a closed-loop system with real-time control is necessary. A critical first step is developing an automated fault detection system that uses sensors to identify printing errors as they occur, which will provide the foundation for a corrective mechanism.

Research has explored several methodologies for detecting extrusion defects, including the analysis of image data[19,20,22–36], temperature [37,38], acoustic emissions[39–42], and electrical resistance[18]. Among these, image analysis using single or multi-camera setups has proven particularly effective[43,26]. For instance, studies have employed Convolutional Neural Networks (CNNs) for analysing images to detect extrusion errors and interlayer imperfections like warping[19,20,22–24,28]. Other approaches use histogram differences to identify defects by comparing prints to reference renders[25–27].

Despite these advancements, significant gaps remain. Studies focus on detecting only specific defects, failing to offer integrated systems that correct multiple error types in real-time[19,22]. Many existing systems are not truly in situ and introduce significant delays. For example, adding up to 50% to the print time, while others do not allow the monitoring of complex geometries[25,26]. In contrast, CNN-based approaches offer a path to feasible real-time detection[44]. However, even these can be limited by static, single-camera setups [28]. The literature thus points toward a need for multi-camera systems for more robust, in-situ error detection and correction.

Addressing these gaps, this study proposes an integrated fault detection system for FFF printers that utilises a multi-camera setup and CNN models. The objective is to create a rapid, accurate, and in-situ system capable of identifying a wide range of common printing errors in real-time. To establish a controlled baseline for validating the core multi-camera fusion concept, the experimental variables were intentionally constrained to a single filament colour, a fixed lighting environment, and a single printer setup. The results demonstrate the system's effectiveness under these constrained experimental conditions: the CNN models achieved a validation accuracy of 97.5% in detecting 12 distinct error types, with model inference times below 220 milliseconds. This research provides a foundation for future work, including the integration of corrective mechanisms and object detection, to ultimately advance the reliability and efficiency of extrusion printing.

Methodology

Hardware setup

The data collection system comprises three cameras: one nozzle camera and two bed cameras. The nozzle camera is mounted on the printer's heat block using a custom-designed bracket for close-up extrusion monitoring. This camera is part of the 3DO nozzle camera kit. The bed cameras detect errors on the sides of the print. These cameras are attached to the printer bed frame, using brackets adapted from an existing open-source design. Each side camera mount features an adjustable swing arm, allowing fine-tuned camera positioning to obtain the optimal field of view for each print. The bed cameras are Raspberry Pi HQ cameras, equipped with CS Mount 6mm 3MP-HD lenses. The mounting of the nozzle camera and side cameras on the printer is shown in **SI Figure 1**. The three cameras capture real-time images, which are transmitted to a hard drive, a Seagate Expansion Portable 2TB hard drive. The side camera was configured to stream at 10 fps at 4056 × 3040, while the nozzle camera streamed at 60 fps at 640 × 480, with exposure set to auto. To manage data collection and printer control, two Raspberry Pi units are implemented: a Raspberry Pi 3 running OctoPrint for remote control and a Raspberry Pi 5, with a 2.4GHz quad-core 64-bit Arm Cortex-A76 CPU and a VideoCore VII GPU, dedicated to storing images. This multi-camera setup is applied to a Prusa MK3S printer, utilising black glass fibre-reinforced polypropylene (PP GF30) as material.

Imaging strategy

The nozzle camera provides a fixed perspective due to the stationary position of the nozzle relative to the camera. This consistency allows for uniform manual cropping of images to exclude visible portions of the heat block and artefacts outside the print plate region, reducing the input size from 640 × 480 to 440 × 320 pixels. These cropped images are then analysed by the dedicated CNN model to detect errors specific to this perspective. For the side cameras, cropping requires manual adjustments. This is because new prints can be positioned differently on the bed, resulting in no fixed reference point for the camera. During the first loop, the operator selects a region of 1600 × 1000 pixels, from images of 4056 × 3040 pixels, focusing on the print area. This region is selected based on the size of the print and its position on the print bed. This manual input for the cropping region is necessary because print sizes and bed positions vary arbitrarily across different jobs, requiring custom adjustment for each print. This cropping region is fixed for all following iterations, as the position of the print does not move with respect to the side cameras, during the same print. The cropped side images are analysed by the respective CNN model, tailored to detect errors observable from the side perspectives, resulting in the classification outputs. The outputs from the three cameras are then passed to the logic system for integration, ensuring prioritised and accurate error reporting.

While the nozzle camera closely monitors the extrusion process, the two bed cameras capture anomalies on the sides of the print in Fig. 1. These three cameras operate in parallel to collect a synchronised, multi-scale dataset covering a comprehensive array of extrusion printing error types.

The development of a fault detection system necessitates a thorough classification of extrusion printing errors encompassing extrusion defects, detail and surface finish anomalies, support and structural failures, layer misalignments, motion faults, and adhesion breakdowns. Extensive datasets for each error category were created by reverse-engineering optimal print parameters to deliberately induce faults, slicing adjustments, and, in some cases, direct printer parameter manipulation like increasing/decreasing the nozzle temperature during printing. A reference dataset of defect-free prints was produced using standard parameters, with representative nozzle- and side-camera images. The datasets were constructed by acquiring images sequentially throughout the printing process by sampling individual frames. Each captured image was manually labelled by assigning the corresponding error class, including “good”, “not printing” and “unclear” in the case where the image could not be used to classify “good” or a defect class.

As shown in Fig. 1, images of twelve error types were collected through the nozzle and side cameras, corresponding to the identification of seven and eight potential errors from the nozzle and side cameras, respectively. These errors were categorised into groups based on their type and root cause. For example, extrusion errors, including over- and under-extrusion, result from incorrect material flow, leading to uneven print line width and thickness or gaps between print layers. Pillowing is caused by inadequate cooling of the top layer, resulting in bumps, and by foreign bodies, which can originate from the nozzle or print bed. Further geometric errors can occur from weak infill, where internal structure of the print is under-extruded, or poor bridging, which occurs when unsupported spans cannot be generated due to the low melt strength of the extruded polymer.

Temperature and pressure-induced errors include stringing, characterised by thin strands of filament between parts of the print due to higher temperature or pressure, and overhang sagging, where unsupported regions creep in their melt state. Next motion and alignment errors include shifted layers, caused by belt slipping or missed motor steps. These failures can also compound. For example, spaghetti failure can occur from poor adhesion propagation or stringing, resulting in

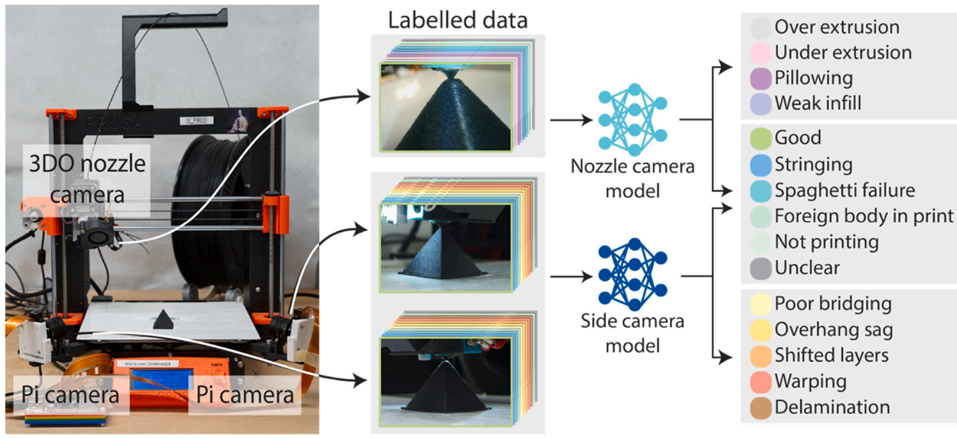

Fig. 1. The data collection and labelling workflow for training the nozzle and side camera models, including the defect types identified by each.

interlaminar failure between print lines, which then results in chaotic filament strands being extruded into free space. Lastly, adhesion errors, such as warping and delamination, are caused by thermal effects such as excessive thermal energy in the part or by insufficient layer bonding, leading to parts lifting from the bed or separating between layers, respectively. Table 1 summarises these errors and their corresponding detection cameras.

The system does not cover all error types. Three errors were found to trigger other detectable faults rather than directly presenting distinct visual features. Support failure and bed separation resulted in spaghetti failure, while overheating led to stringing, overhang sagging, poor bridging and warping. A total of seven errors were excluded from the developed detection system either because they require additional comparison to a digital twin or they are not visible. Small features not printing, poor surface quality at support regions, missing layers and first layer problems were omitted as they are not easily visible without a comparison to external references, such as CAD files. Further errors, such as ripples and echoes, infill showing protruding from the exterior of the print and elephant foot, were also excluded due to their unreliability in being detected visually with the proposed camera system. Ripples and echoes could potentially be addressed by vibration or accelerometer sensors, infill protrusion by laser scanning or force sensors, and elephant foot by force sensors or thermal imaging.

Dataset and ML model description

Around 90,000 images were collected and labelled to generate a dataset which covered the 3D printing errors introduced above. The total number of nozzle camera and side camera images per class in the useful dataset is shown in **SI Table 1**. To ensure geometric variability, a diverse range of parts was printed. As shown in **SI Figure 2**, these included basic shapes like cubes, cylinders, etc., functional test parts like bridges, overhang tests, and more complex forms like F-shapes and T-shapes. To generate sufficient data for each defect class, images were acquired sequentially throughout the printing process by sampling individual frames while intentionally failing prints through modified process parameters or manual intervention.

SI Table 2 specifies the multiple, distinct geometries used to generate data for each specific error, ensuring that defects were captured across a variety of shapes and conditions. For example, "warping" was induced on low cubes, pyramids, and A-shape blocks, among others. Similarly, "poor bridging" was captured on multiple bridge tests and H-shaped parts. The images were labelled according to their respective error classes and subsequently balanced. For each error type, a total of 2,000 images from both side cameras combined, and 2,400 images from the nozzle model, respectively, were selected from the total image dataset for class balancing before training the model.

Both the nozzle and side camera models use the pretrained You Only Look Once version 8 (YOLOv8) image classification model. To enhance the model's generalisation capability, the YOLO model's default data augmentation techniques, such as horizontal flipping, rotations and colour adjustments, such as altering brightness, contrast and saturation, were applied to the training data within the model framework. To evaluate generalisation to unseen geometries, an external dataset featuring a distinct shape shown in **SI Figure 3** was created as a baseline demonstration under controlled conditions. This validates error detection capabilities on unseen structures. The external dataset prints were created under identical experimental conditions, including the same filament colour, fixed ambient lighting, and printer setup. In future work, the external test set will include multiple geometries to comprehensively assess robustness.

The YOLOv8s classification model is chosen for its superior accuracy and efficiency over architectures such as ResNet50 and VGG16 on a labelled 3D-printing dataset to discriminate between defect-free prints and various error classes in real time[46,47]. The YOLOv8 architecture shown in Fig. 2 comprises Convolution Blocks (2D convolution with 3×3 kernel size, 2×2 stride and 1×1 zero padding followed by batch normalization and SiLU activation), C2f blocks (cross-stage partial connections with bottleneck layers for lightweight yet expressive feature reuse), and a Classify Block (convolution, max-pooling, dropout, and fully connected layer) that together balance computational efficiency with high classification performance. Model training involves adjusting the final classification head on our error-labelled images using standard transfer-learning workflows in Ultralytics' YOLOv8 framework. The hyperparameters used for model training are shown in **SI Table 3**. This is followed by deployment for live camera feeds to flag print faults as they occur.

Table 1
Error types and cameras used for detection.

Error group	Error	Detecting camera
Extrusion errors	Over-extrusion	Nozzle camera
	Under-extrusion	Nozzle camera
Detail issues	Pillowing	Nozzle camera
	Foreign body in print	Nozzle and side cameras
Structural errors	Weak infill	Nozzle camera
	Poor bridging	Side cameras
Temperature errors	Stringing	Nozzle and side cameras
	Overhang sagging	Side cameras
Alignment and motion errors	Shifted Layers	Side cameras
	Spaghetti failure	Nozzle and side cameras
Adhesion errors	Warping	Side cameras
	Delamination	Side cameras

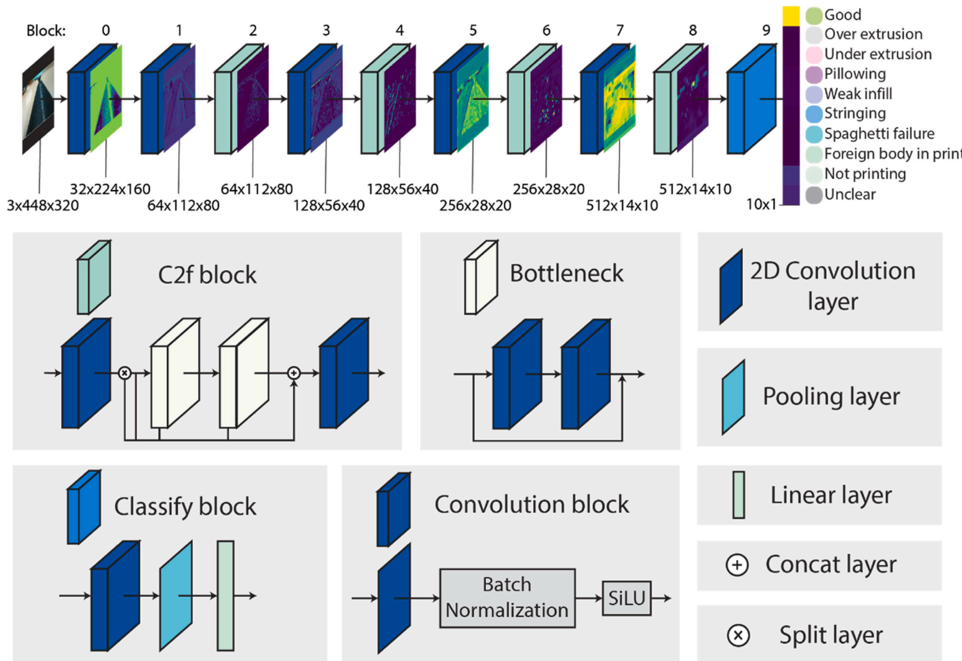

Fig. 2. An overview of the YOLOv8s classification model[45]. The figure illustrates the model's overall architecture, provides a detailed expansion of its main computational blocks, and shows the resulting activation maps when applied to a sample nozzle camera image. The same model architecture was trained on the image dataset sourced from both side cameras to develop the side camera model.

Results and discussion

The nozzle camera model is trained for 10 epochs using a batch size of 16, an input resolution of 448 × 320 pixels, and a learning rate of 1.5 × 10^{-3}. The side camera model is trained for 50 epochs with a batch size of 23, an input resolution of 512 × 512 pixels after rescaling, and a learning rate of 1.1 × 10^{-3}. The best-performing weights for the side camera model, corresponding to the highest validation accuracy, are obtained at 30 epochs. Relatively low number of training epochs in both cases is due to starting the training from a pre-trained model. The cross-entropy loss curves along with the validation accuracy curves for both models are shown in the **SI Figure 4**. Both models show strong performance as seen from the cross-entropy loss, accuracy and F1 score values shown in Table 2. Cross-entropy loss quantifies how well predicted probabilities align with true labels, accuracy shows the proportion of correct predictions, and the F1 score balances precision (measure of lack of false positives) and recall (measure of true positives) for a comprehensive evaluation of classification performance.

The model uses the training data to optimise the network's weights during training and makes predictions on the validation dataset at each training step to ensure that the model does not overfit the training dataset. Confusion matrices are used to analyse how effectively the model distinguishes between different classes, while Grad-CAM heatmaps[48] are used to provide visual insights into which input regions most influence the model's predictions. This helps in interpreting the model's decision-making patterns.

The confusion matrices are computed for the training, validation, test and external test datasets. Each confusion matrix represents the true classification along the horizontal axis and the model's prediction along the vertical axis. The confusion matrices, shown in Fig. 3**a**, were computed to visualise the accuracy of the model's prediction across different model development classes. The points of intersection of the actual and predicted classes, as the matrix diagonal in blue, represent the correct classifications, and all points outside the diagonal represent misclassifications. The high values along the diagonals of confusion matrices for these two datasets in both models indicate that the models train with high accuracy on the training data and generalise well to the validation data. The test and external test datasets are not available to the model during training, and predictions are made on these datasets after the model has been trained. The confusion matrices of the test dataset for both models show that they generalise well to these datasets, with only a few misclassifications. Though fairly good performance is achieved in the external test dataset, there are more misclassifications when compared to the internal test dataset. This has been investigated further using Gradient-weighted Class Activation Mapping (Grad-CAM) heatmaps. Grad-CAM heat maps are used to analyse the classifications of the model on the test and external test datasets. This helps ensure that the model has learned robust features and can also highlight potential weaknesses or biases in the model's decision-making process.

The heat maps visualise the importance of image regions for a CNN's decision by computing the gradients of the target class score with respect to the feature maps of a convolutional layer, then weighting and summing those maps to produce a localization mask. These are typically obtained from the final convolutional layers of the CNN network. They can be superimposed on the original image to identify the features of the image that are responsible for the network's classification decision. This

Table 2
Nozzle camera model and side camera model performance metric.

Metric	Training		Validation		Test		External test	
	Nozzle	Side	Nozzle	Side	Nozzle	Side	Nozzle	Side
Accuracy (%)	99.11	99.81	97.68	97.61	97.09	97.86	93.18	89.3
Loss	0.07	0.04	0.07	0.12	0.07	0.09	0.27	0.39
F1 score (%)	99.11	99.82	97.66	97.64	97.1	97.93	91.09	84.06

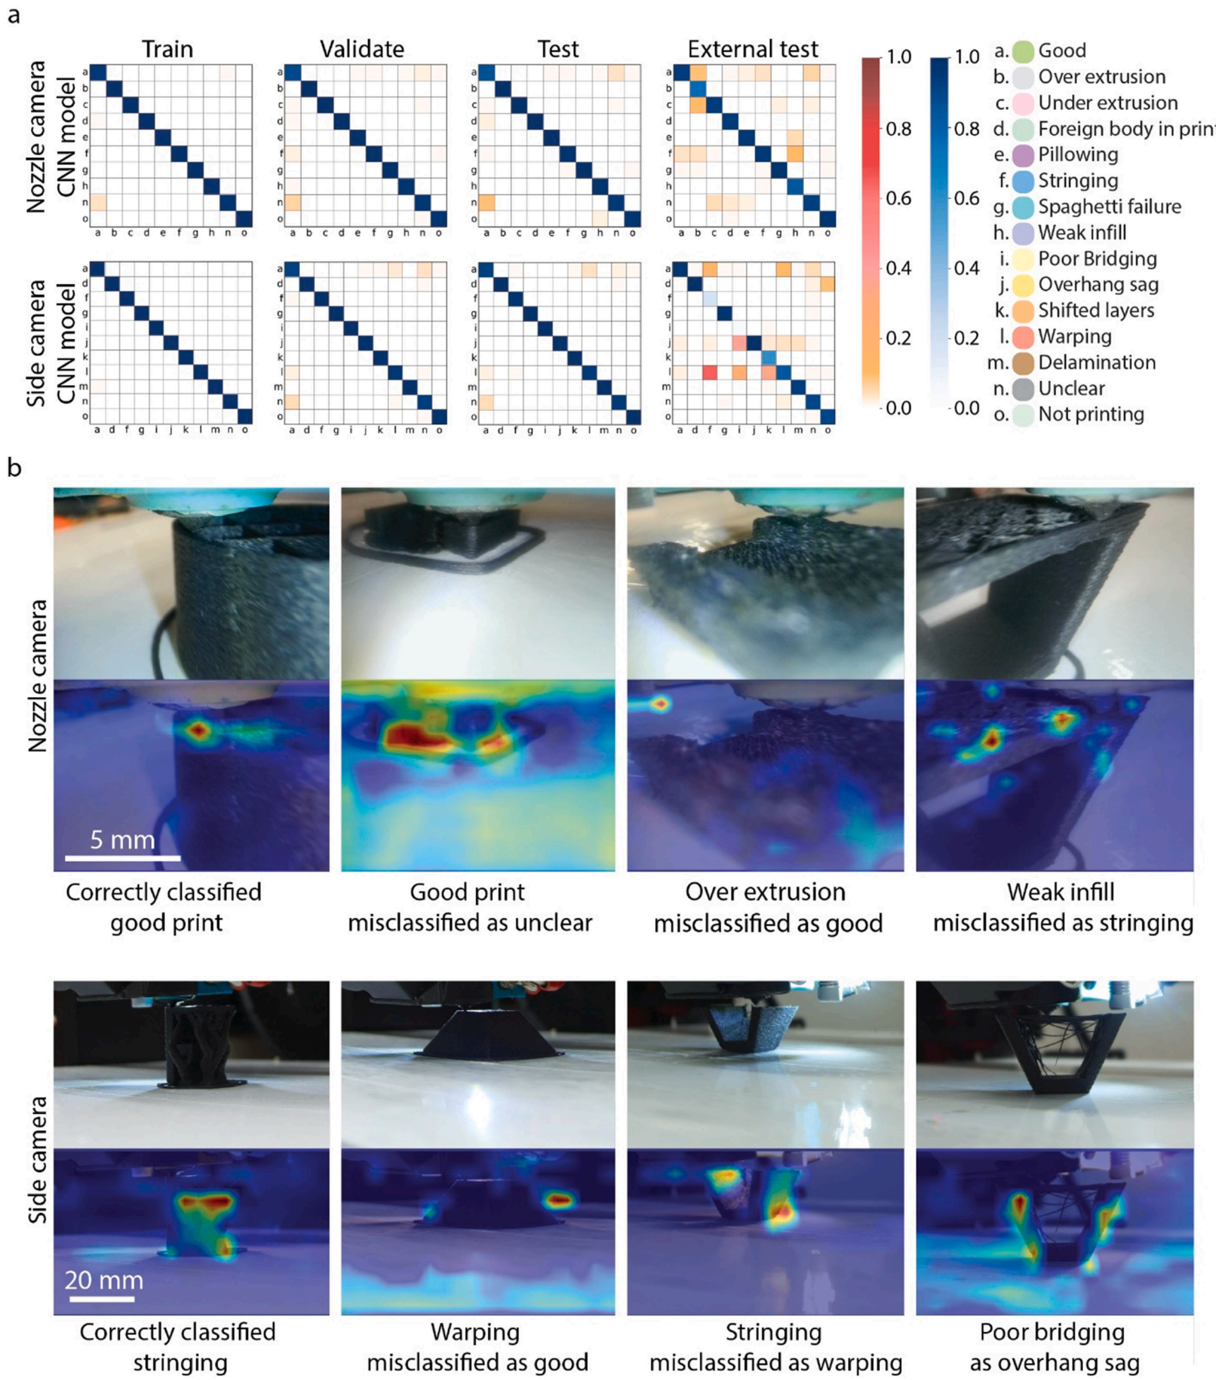

Fig. 3. Confusion matrices and Grad-CAM heatmaps from the CNN model. a. The confusion matrices for training, validation and test datasets for the nozzle camera and side camera models. The blue colour bar represents correct classifications, while the red colour bar represents misclassifications. b. Representative Grad-CAM heatmaps of the predicted class at the output of the 8th convolution block in the YOLOv8 model's backbone, chosen for its optimal balance of semantic abstraction and spatial detail.

heat map highlights areas with high positive influence on the class prediction, effectively attributing spatial relevance to the network's output. From the image masks, the model attends to relevant regions for correctly classified images, while for misclassified images, the heatmaps can reveal which image areas may have led to incorrect decisions.

Nozzle camera model performance

Fig. 3**b** shows representative example images, selected to illustrate typical cases for both the nozzle and side camera models. For the nozzle camera model, an example is shown where a good print is identified correctly. As can be seen in the overlay, the region of highest activation corresponds directly to the area around the nozzle. The same is seen in correctly classified images of under-extrusion and over-extrusion (see **SI Figure 5**).

It is seen from the test dataset's confusion matrix that around 9% of good print images are misclassified as "unclear" for the nozzle model, as shown in the example image in Fig. 3**b**. Investigating this with Grad-CAM heatmaps shows that the model does not focus on the nozzle but on other irrelevant parts of the image, resulting in misclassified images. A representative example of this behaviour is included in Fig. 3**b**. These misclassifications occur because the "unclear" class was defined subjectively, without clearly distinguishable features. However, these misclassifications have no significant implications for the detection system. In addition, poor lighting results in poor distinction between the nozzle tip and the printed part. In the external test dataset, 9% of over-extrusion images are classified as “good” or under-extrusion. As visualised by the corresponding Grad-CAM overlay in Fig. 3**b**, the model does not focus on the nozzle tip in this case either, since the image is blurry or there is overexposure in the image. These observations demonstrate that optimum lighting on the nozzle tip is essential to recognise good, over-extrusion and under-extrusion errors during the printing process. Adaptation to filaments with different colours and reflectivity may require different lighting conditions to improve reliable defect recognition.

For the classification of further errors by the nozzle camera model, the model focuses on the defect features on the print body rather than the nozzle tip to make correct predictions. 11% of weak infill images are classified as stringing by the model in the external test dataset. Corresponding Grad-CAM heat maps show that for these misclassified images, some string-like features are present on the print layer along with the weak infill, and the model seems to focus on this for classifying the image.

Side camera model performance

When studying the side camera model, the external test set confusion matrix shows that images with stringing, poor bridging and shifted layer defects are the most misclassified. For a representative example of correctly classified stringing in the internal test dataset, as visualised in Fig. 3**b**, the model focuses on the stringing features in the printed structure. At some angles, string-like features are not visible to the side camera, and hence, the model is unable to identify them. This explains why 15% of stringing defects are classified as “good” prints. This also indicates that the nozzle camera will make more reliable predictions for stringing errors.

From the external test set 65%, 24% and 34% of stringing, poor bridging and shifted layer defects, respectively, are classified as warping. Further analysis shows that the images misclassified as warping also exhibited visible signs of warping, which the model correctly identified and used for classification. This applies to the poor bridging and shifted layer cases as well (see **SI Figure 5**). This highlights that the model cannot yet recognise multiple defects present in a single image. Furthermore, it is seen from the confusion matrix that 38% of poor bridging images are misclassified as overhang sagging in the external test set. The Grad-CAM heatmap for a representative example, shown in Fig. 3**b**, reveals that both overhang sagging and poor bridging defects are present in the print, and the model focuses on the overhang sag features. This highlights the importance of detecting multiple errors in a single print and the potential to use object recognition models instead of classification models in future work.

From the external test set, 11% of warping errors are misclassified as “good” prints. As shown by the corresponding Grad-CAM heatmap in Fig. 3**b** the model does not focus on the warped edges, which usually is where minor warping is onset. This indicates that the model struggles to detect subtle edge warping. Adopting an object detection approach, i.e., labelling warped edges with bounding boxes, could help the model better learn and localise minimal warping features.

As images were sampled as individual frames from a continuous acquisition process, it is possible that highly correlated frames from the same defect event were distributed across both the training and testing sets. This lack of strict separation could potentially lead to data leakage and an inflated test accuracy. However, this concern is mitigated by two key observations. First, the Grad-CAM heatmaps for most correctly classified images effectively highlight the actual defect regions, indicating that the model is focusing on relevant features rather than spurious correlations. Furthermore, the model's strong performance on the external test dataset, which consisted of an exclusive geometry not seen during training, reinforces this point. The minimal drop in accuracy on this external set suggests that the model learned to generalise beyond specific training examples and demonstrates a meaningful degree of robustness. For a more rigorous evaluation, subsequent work should employ grouped splits, where all frames from a single print job are kept within the same dataset partition, ensuring that the geometries in the train and test sets are entirely mutually exclusive.

Integrated fault detection system

The detection system operates in a continuous loop, capturing three images per timestep at a sampling rate of 1 Hz. One image is obtained from the nozzle camera, and images are obtained from the left and right-side cameras. These images are processed separately, resulting in three outputs that need to be integrated to reach a single decision as shown in Fig. 4**a**. The initial images are cropped before being passed through the models. Cropping ensures the images better resemble the training data and reduces memory limitations. Additionally, cropping eliminates unnecessary peripheral noise, allowing the analysis to focus on the relevant print areas.

The logic system, in Fig. 4**b**, transforms raw model outputs into actionable conclusions. First the individual answers for each camera are derived, based on consecutive outputs of that camera. These answers are combined into a unified system-level conclusion. The two-step process ensures consistency and prioritisation. Each camera’s outputs are saved in separate lists. The output lists undergo a filtering process to address cases where outputs such as “good”, “unclear” or “not printing” appear. These require special handling to avoid disrupting error detection.

First, considering the individual cameras, if a model predicts “good,” this output is tracked, and a counter is incremented for consecutive “good” predictions. Once four consecutive “good” predictions are observed, the list of previous outputs is cleared for that camera. This clearing ensures that, for example, if there are 1,500 predictions, with 1,496 classified as “good” and only 4 scattered as stringing, the system avoids incorrectly returning stringing. Secondly, “unclear” outputs are temporarily included for evaluation but subsequently removed. Therefore, "unclear" images do not disturb the decision-making sequence. Thirdly, outputs “not printing” are excluded unless detected simultaneously by two or more cameras, as a single “not printing” prediction could result from nozzle movements rather than an actual fault.

Error confirmation relies on consistency across iterations. If an output, such as “stringing”, is repeated consistently across four or more iterations for one camera, it is parsed as the camera’s answer. This is also the case when two cameras detect the same error simultaneously twice in a row. Similarly, when all three cameras agree on the same error, this answer will be returned. In all three cases, the system tolerates intermittent “good” or "unclear" predictions within the sequence to account for model inaccuracies. This logic is statistically robust, ensuring a minimum detection accuracy of 99% even when individual model accuracies are approximately 97%.

The final step in the fault detection system integrates the outputs from the three cameras into a single cohesive conclusion. This requires an integration logic that prioritises errors based on context as shown in Table 3. For instance, when side cameras detect overhang sagging, poor bridging or delamination, the nozzle camera’s output is ignored. If the side cameras do not detect any of these three errors and the nozzle camera identifies over-extrusion or under-extrusion, the side camera outputs are disregarded for that iteration. Disregarding cameras in these scenarios ensures that errors detectable by one camera, but not by others, do not introduce incorrect conclusions due to overlapping features. So, by prioritising and ignoring specific camera outputs, the system avoids conflicting conclusions. The inference times from the individual CNN models were observed to be approximately 220 milliseconds. System latency was evaluated by separating computational processing from hardware capture overhead. Processing latency per iteration, measured by replaying 224 image sets through the inference

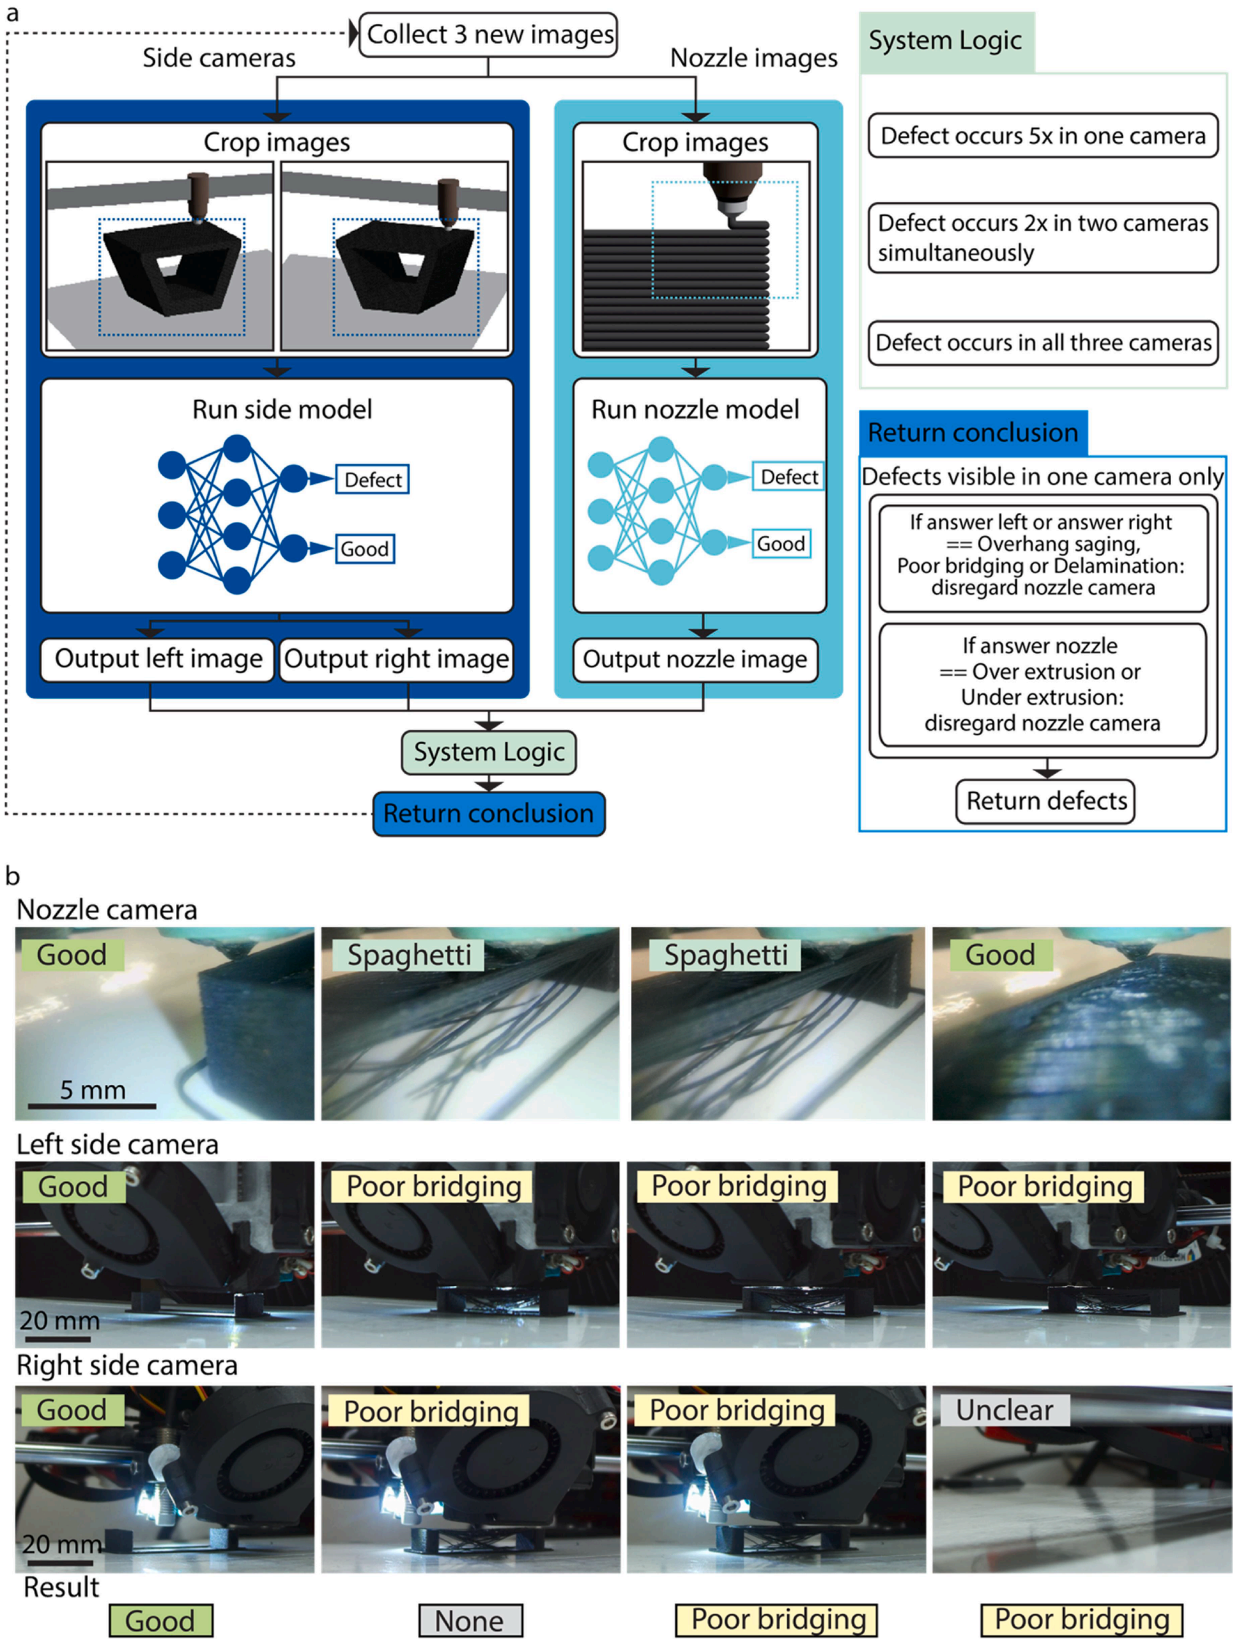

Fig. 4. The operational workflow of the multi-camera fault detection system. a. A schematic illustrating the system logic, where images from three cameras are individually cropped, classified by CNN models, and then integrated. The system logic applies rules for error confirmation based on detection consistency and uses a prioritization system to resolve conflicting outputs from different camera perspectives. b. A demonstration of the system's performance over four iterations. Iteration 1: "Good" (all cameras "good"). Iteration 2: "None" (intermittent state - 2 side cameras "Poor Bridging," nozzle "Spaghetti"). Iteration 3: "Poor Bridging" confirmed (2 cameras agree for 2 consecutive frames). Iteration 4: "Poor Bridging" maintained. Prioritisation logic requires two cameras to agree on the same error across two consecutive frames, ignoring conflicting single-camera predictions.

Table 3
Priority rules for defect detection, where defects in side cameras (overhang sagging, poor bridging, delamination) or nozzle camera (over/under-extrusion) take precedence.

Condition	Action
Left or right camera: Overhang sagging, Poor bridging, or Delamination	Disregard the nozzle camera, return left and right camera prediction
Nozzle camera: Over-extrusion or Under-extrusion	Disregard left and right camera, return nozzle camera prediction

pipeline on a laptop workstation equipped with an Intel Core i7-1365U CPU at 1.80 GHz and 16GB RAM, averaged 343 milliseconds. The capture latency for the two side cameras is estimated at approximately 2.30 seconds because the script initialises each camera sequentially and writes images to disk via Picamera2. In contrast, the nozzle camera capture is much faster, estimated at 0.10 seconds, as it reads directly from a continuous video stream via OpenCV. Therefore, the total system latency is approximately 2.40 seconds plus the processing time, with the majority of the delay caused by the camera initialisation method. Switching to continuous capture streams for all cameras could reduce this hardware delay.

This integration successfully establishes a baseline for real-time, multi-camera fault detection. An example of the system's operation over four iterations is shown in Fig. 4**b**. Our model successfully classifies 12 distinct 3D printing defects with 97% accuracy under the constrained setup with single filament colour and fixed lighting. In future work, more data will be generated in diverse conditions with varying filament colours, lighting, and printer setups and cross-validated to ensure robust performance. Most published systems identify fewer defects, typically between four and six types[30,33,34,36]. Our work is notable because it achieves this broad classification scope while minimising disruptions caused by transient or isolated mispredictions.

The current system utilises a single-label classification approach. This can be extended by using object detection models, which are capable of multi-label classification. In such a system, the decision algorithm would be modified to process multiple simultaneous defect detections and their severities to arrive at a reliable conclusion. Another limitation is our model's current specialization for black-coloured prints. A crucial area for future work will be to quantify the model's generalizability by testing it against a variety of material colours and types, fluctuating ambient lighting, and different printer models. These investigations will be essential for transitioning the system from a proof-of-concept to a universally applicable tool.

Conclusions

The feasibility and efficacy of a real-time, multi-camera fault detection system marks a substantial advance over prior single-view, single-defect approaches. By integrating two side-view cameras with a nozzle-mounted camera and employing two specialised convolutional neural network (CNN) classifiers, it is shown that a diverse set of twelve error modes, including over- and under-extrusion, stringing, warping, and delamination, can be detected under constrained conditions of single filament colour, single lighting setup and single printer configuration. These results compare favourably to earlier studies that reported lower detection rates for only one or two defect types under constrained imaging conditions. A broader generalisation of this model remains a goal for future work.

The consequences of these findings are twofold. First, in situ error identification without print interruption can drastically reduce material waste, energy consumption, and manufacturing time by enabling early intervention or automatic corrective actions. Second, our logic-based fusion framework, which reconciles potentially conflicting outputs from multiple cameras, enhances robustness and paves the way for adaptive control strategies in additive manufacturing. These advantages are substantiated through quantitative performance metrics and representative case studies highlighting scenarios in which multi-angle imaging corrected misclassifications that single-view systems would have missed.

Beyond desktop-scale FFF, the methodology can generalise to other extrusion technologies and larger industrial platforms, suggesting broad applicability in smart manufacturing. By providing an open-ended architecture for real-time monitoring and classification, this study lays the groundwork for future integration of closed-loop corrective mechanisms and contributes a critical step toward fully autonomous, fault-tolerant 3D printing.

Declaration of AI

During the preparation of this work, the authors used OPEN AI ChatGPT 4.0 to check for spelling and grammar mistakes and improve readability. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the published article.

CRediT authorship contribution statement

Shanthalakshmi Kilambi: Writing – review & editing, Writing – original draft, Visualization, Validation, Supervision, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Aster Tournoy:** Writing – review & editing, Writing – original draft, Software, Visualization, Validation, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Muhamad Amani:** Writing – review & editing, Supervision, Investigation, Formal analysis, Conceptualization. **Jovana Jovanova:** Writing – review & editing, Supervision, Conceptualization. **Baris Caglar:** Writing – review & editing, Validation, Supervision, Conceptualization. **Kunal Masania:** Writing – review & editing, Writing – original draft, Visualization, Validation, Supervision, Investigation, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

We acknowledge funding from the National Growth Fund Program NXTGEN HIGHTECH, the European Union (ERC Consolidator Grant, AM-IMATE, 101088968) and support of Tanaruz B.V.

Supplementary materials

Supplementary material associated with this article can be found, in the online version, at doi:10.1016/j.addlet.2026.100360.

Data availability

The data that support the findings of this study are available from the corresponding author upon reasonable request. Source codes will be made available through our GitLab page.

References

[1] L. Rettenberger, N. Beyer, I. Sieber, M. Reischl, Fault detection in 3D-printing with deep learning, in: 2024 IEEE International Conference on Consumer Electronics (ICCE), 2024, pp. 1–4, https://doi.org/10.1109/ICCE59016.2024.10444198.

[2] Z. Jin, Z. Zhang, K. Demir, G.X. Gu, Machine learning for advanced additive manufacturing, Matter 3 (5) (2020) 1541–1556, https://doi.org/10.1016/j.matt.2020.08.023.

[3] J. Patadiya, A. Gawande, G. Joshi, B. Kandasubramanian, Additive manufacturing of shape memory polymer composites for futuristic technology, Ind. Eng. Chem. Res. 60 (Oct. 2021), https://doi.org/10.1021/acs.iecr.1c03083.
[4] J.Z. Gul, et al., 3D printing for soft robotics - a review, Sci. Technol. Adv. Mater. 19 (1) (2018) 243–262, https://doi.org/10.1080/14686996.2018.1431862.
[5] A. Ahmed, S. Arya, V. Gupta, H. Furukawa, A. Khosla, 4D printing: fundamentals, materials, applications and challenges, Polymer 228 (2021) 123926, https://doi.org/10.1016/j.polymer.2021.123926.
[6] A. Behera, Smart materials based additive manufacturing, in: S. Rajendrachari (Ed.), Practical Implementations of Additive Manufacturing Technologies, Springer Nature, Singapore, 2024, pp. 153–175, https://doi.org/10.1007/978-981-99-5949-5_7.
[7] M.Y. Khalid, Z.U. Arif, A. Tariq, M. Hossain, K.Ahmed Khan, R. Umer, 3D printing of magneto-active smart materials for advanced actuators and soft robotics applications, Eur. Polym. J. 205 (2024) 112718, https://doi.org/10.1016/j.eurpolymj.2023.112718.
[8] R.T. Shafranek, S.C. Millik, P.T. Smith, C.-U. Lee, A.J. Boydston, A. Nelson, Stimuli-responsive materials in additive manufacturing, Prog. Polym. Sci. 93 (2019) 36–67, https://doi.org/10.1016/j.progpolymsci.2019.03.002.
[9] X. Zhou, et al., Advances in 3D/4D printing of mechanical metamaterials: from manufacturing to applications, Compos. B 254 (2023) 110585, https://doi.org/10.1016/j.compositesb.2023.110585.
[10] P. Rastogi, B. Kandasubramanian, Breakthrough in the printing tactics for stimuli-responsive materials: 4D printing, Chem. Eng. J. 366 (2019) 264–304, https://doi.org/10.1016/j.cej.2019.02.085.
[11] A. Megdich, H. Mohamed, L. Laperrière, A review on 4D printing: material structures, stimuli and additive manufacturing techniques, Mater. Lett. 337 (Feb. 2023) 133977, https://doi.org/10.1016/j.matlet.2023.133977.
[12] D. Song, A.M. Chung Baek, J. Koo, M. Busogi, N. Kim, Forecasting warping deformation using multivariate thermal time series and k-nearest neighbors in fused deposition modeling, Appl. Sci. 10 (24) (2020) 8951, https://doi.org/10.3390/app10248951. Art.
[13] A. Vashisth, R.E. Healey, M.J. Pospisil, J.H. Oh, M.J. Green, Continuous processing of pre-pregs using radio frequency heating, Compos. Sci. Technol. 195 (Jul. 2020) 108211, https://doi.org/10.1016/j.compscitech.2020.108211.
[14] C. Houriet, B. Ulyanov, J.-A. Pascoe, K. Masania, Wood-inspired interlocking junctions using 3D-printed liquid crystal polymers, Addit. Manuf. 97 (Jan. 2025) 104590, https://doi.org/10.1016/j.addma.2024.104590.
[15] S.K. Ammu, et al., 3D printing of lead-free piezoelectric ultrasound transducers, Adv. Mater. Technol. 9 (23) (2024) 2400858, https://doi.org/10.1002/admt.202400858.
[16] C. Wang, Z. Zhao, Z. Li, Y. Jia, A.A. Sharma, X.S. Zhang, Direct ink writing of magnetic soft materials with optimized printing path, Addit. Manuf. 105 (May 2025) 104770, https://doi.org/10.1016/j.addma.2025.104770.
[17] I.H.M.S. Nettersheim, N.S.G. Sotelo, J.C. Verdonk, K. Masania, Engineered living composite materials, Compos. Sci. Technol. 256 (Sep. 2024) 110758, https://doi.org/10.1016/j.compscitech.2024.110758.
[18] H. Jonkers, A. Dijkshoorn, and S. Stramigioli, 'FFF print defect characterization through in-situ electrical resistance monitoring', 2024, doi: 10.21203/rs.3.rs-3956508/v1.
[19] D.A.J. Brion, S.W. Pattinson, Generalisable 3D printing error detection and correction via multi-head neural networks, Nat. Commun. 13 (1) (Dec. 2022), https://doi.org/10.1038/s41467-022-31985-y.
[20] D.A.J. Brion, S.W. Pattinson, Quantitative and real-time control of 3d printing material flow through deep learning, Adv. Intell. Syst. 4 (11) (Nov. 2022), https://doi.org/10.1002/aisy.202200153.
[21] B.T. Wittbrodt, et al., Life-cycle economic analysis of distributed manufacturing with open-source 3-D printers, Mechatronics 23 (6) (2013) 713–726, https://doi.org/10.1016/j.mechatronics.2013.06.002.
[22] Z. Jin, Z. Zhang, G.X. Gu, Automated real-time detection and prediction of interlayer imperfections in additive manufacturing processes using artificial intelligence, Adv. Intell. Syst. 2 (1) (Dec. 2019), https://doi.org/10.1002/aisy.201900130.
[23] Z. Jin, Z. Zhang, J. Ott, G.X. Gu, Precise localization and semantic segmentation detection of printing conditions in fused filament fabrication technologies using machine learning, Addit. Manuf. 37 (2021) 101696, https://doi.org/10.1016/j.addma.2020.101696.
[24] Z. Jin, Z. Zhang, G.X. Gu, Autonomous in-situ correction of fused deposition modeling printers using computer vision and deep learning, Manuf. Lett. 22 (Oct. 2019) 11–15, https://doi.org/10.1016/J.MFGLET.2019.09.005.
[25] A.L. Petsiuk, J.M. Pearce, Open source computer vision-based layer-wise 3D printing analysis, Addit. Manuf. 36 (Dec. 2020) 101473, https://doi.org/10.1016/J.ADDMA.2020.101473.
[26] A. Petsiuk, J.M. Pearce, Towards smart monitored AM: open source in-situ layer-wise 3D printing image anomaly detection using histograms of oriented gradients and a physics-based rendering engine, Addit. Manuf. 52 (Apr. 2022) 102690, https://doi.org/10.1016/J.ADDMA.2022.102690.
[27] S. Nuchitprasitchai, M. Roggemann, J.M. Pearce, Factors effecting real-time optical monitoring of fused filament 3D printing, Prog. Addit. Manuf. 2 (3) (Sep. 2017) 133–149, https://doi.org/10.1007/s40964-017-0027-x.
[28] D.A.J. Brion, M. Shen, S.W. Pattinson, Automated recognition and correction of warp deformation in extrusion additive manufacturing, Addit. Manuf. 56 (Aug. 2022) 102838, https://doi.org/10.1016/J.ADDMA.2022.102838.
[29] A. Saluja, J. Xie, K. Fayazbakhsh, A closed-loop in-process warping detection system for fused filament fabrication using convolutional neural networks, J. Manuf. Process. 58 (2020) 407–415, https://doi.org/10.1016/j.jmapro.2020.08.036.
[30] N.B.A. Karna, M.A.P. Putra, S.M. Rachmawati, M. Abisado, G.A. Sampedro, Toward accurate fused deposition modeling 3D printer fault detection using improved yolov8 with hyperparameter optimization, IEEE Access 11 (2023) 74251–74262, https://doi.org/10.1109/ACCESS.2023.3293056.
[31] Y. Wang, et al., A CNN-based adaptive surface monitoring system for fused deposition modeling, IEEE/ASME Trans. Mechatron. 25 (5) (2020) 2287–2296, https://doi.org/10.1109/TMECH.2020.2996223.
[32] Y. Fu, A.R.J. Downey, L. Yuan, H.-T. Huang, Real-time structural validation for material extrusion additive manufacturing, Addit. Manuf. 65 (2023) 103409, https://doi.org/10.1016/j.addma.2023.103409.
[33] M. Bhattacharya, M. Penica, E. O'Connell, M. Hayes, AI-driven real-time failure detection in additive manufacturing, Procedia Comput. Sci. 232 (2024) 3229–3238, https://doi.org/10.1016/j.procs.2024.02.138.
[34] Z. Zhang, I. Fidan, M. Allen, Detection of material extrusion in-process failures via deep learning, Inventions 5 (3) (Sep. 2020) 1–9, https://doi.org/10.3390/inventions5030025.
[35] I.A. Okaro, S. Jayasinghe, C. Sutcliffe, K. Black, P. Paoletti, P.L. Green, Automatic fault detection for laser powder-bed fusion using semi-supervised machine learning, Addit. Manuf. 27 (2019) 42–53, https://doi.org/10.1016/j.addma.2019.01.006.
[36] H. Kim, H. Lee, J.-S. Kim, S.-H. Ahn, Image-based failure detection for material extrusion process using a convolutional neural network, Int. J. Adv. Manuf. Technol. 111 (5) (Nov. 2020) 1291–1302, https://doi.org/10.1007/s00170-020-06201-0.
[37] E. Ferraris, J. Zhang, B.V. Hooreweder, Thermography based in-process monitoring of fused filament fabrication of polymeric parts, CIRP Ann. 68 (1) (Jan. 2019) 213–216, https://doi.org/10.1016/J.CIRP.2019.04.123.
[38] H. Baumgartl, J. Tomas, R. Buettner, M. Merkel, A deep learning-based model for defect detection in laser-powder bed fusion using in-situ thermographic monitoring, Prog. Addit. Manuf. (Sep. 2020), https://doi.org/10.1007/s40964-019-00108-3.
[39] H. Wu, Z. Yu, Y. Wang, A new approach for online monitoring of additive manufacturing based on acoustic emission, in: International Manufacturing Science and Engineering Conference 3, Joint MSEC-NAMRC Symposia, 2016 V003T08A013, https://doi.org/10.1115/MSEC2016-8551 vol. Volume 3: Joint MSEC-NAMRC Symposia.
[40] H. Wu, Z. Yu, Y. Wang, Real-time FDM machine condition monitoring and diagnosis based on acoustic emission and hidden semi-Markov model, Int. J. Adv. Manuf. Technol. 90 (5) (May 2017) 2027–2036, https://doi.org/10.1007/s00170-016-9548-6.
[41] P. Becker, C. Roth, A. Roennau, R. Dillmann, Acoustic anomaly detection in additive manufacturing with long short-term memory neural networks, in: 2020 IEEE 7th International Conference on Industrial Engineering and Applications (ICIEA), 2020, pp. 921–926, https://doi.org/10.1109/ICIEA49774.2020.9102002.
[42] W.J. Staszewski, B.C. Lee, L. Mallet, F. Scarpa, Structural health monitoring using scanning laser vibrometry: I. Lamb wave sensing, Smart Mater. Struct. 13 (2) (Feb. 2004) 251, https://doi.org/10.1088/0964-1726/13/2/002.
[43] A. Oleff, B. Küster, M. Stonis, L. Overmeyer, Process monitoring for material extrusion additive manufacturing: a state-of-the-art review, Prog. Addit. Manuf. 6 (4) (Dec. 2021) 705–730, https://doi.org/10.1007/s40964-021-00192-4.
[44] O. Fink, Q. Wang, M. Svensén, P. Dersin, W.-J. Lee, M. Ducoffe, Potential, challenges and future directions for deep learning in prognostics and health management applications, Eng. Appl. Artif. Intell. 92 (Jun. 2020) 103678, https://doi.org/10.1016/j.engappai.2020.103678.
[45] P. Hidayatullah, N. Syakrani, M. R. Sholahuddin, T. Gelar, and R. Tubagus, 'YOLOv8 to YOLO11: a comprehensive architecture in-depth comparative review', Apr. 08, 2025, *arXiv*: arXiv:2501.13400. doi: 10.48550/arXiv.2501.13400.
[46] Y. Chen, The investigation of performance comparison for VGG, YOLO, and DINO in image classification, Highlights Sci. Eng. Technol. 85 (Mar. 2024) 984–990, https://doi.org/10.54097/9bgem219.
[47] V. Patel, M. Kanojia, V. Nair, Exploring the potential of ResNet50 and yolov8 in improving breast cancer diagnosis: a deep learning perspective, Int. J. Comput. Inf. Syst. Ind. Manag. Appl. 16 (3) (Jul. 2024) 16.
[48] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, 'Grad-CAM: visual explanations from deep networks via gradient-based localization', Dec. 03, 2019. doi: 10.1007/s11263-019-01228-7.