Measurement Science
and Technology

PAPER • **OPEN ACCESS**

High-speed multimodal actionable porosity prediction for enabling in-process control of selective laser sintering

To cite this article: Mathieu Vandecasteele *et al* 2026 *Meas. Sci. Technol.* **37** 025602

View the article online for updates and enhancements.

You may also like

- Sub-millisecond keyhole pore detection in laser powder bed fusion using sound and light sensors and machine learning
Zhongshu Ren, Jiayun Shao, Haolin Liu et al.
- Real-time defect detection of laser additive manufacturing based on support vector machine
Tao Liu, Lei Huang and Bo Chen
- Predicting surface roughness of overhang structures by registering and fusing sensing signals and morphology data in LPBF
Longchao Cao, Tao Zhang, Xufeng Huang et al.



This content was downloaded from IP address 146.103.254.11 on 19/11/2026 at 12:41

IOP Publishing Meas. Sci. Technol. 37 (2026) 025602 https://doi.org/10.1088/1361-6501/ae2cbe

Measurement Science and Technology

Check for updates

PAPER

High-speed multimodal actionable porosity prediction for enabling in-process control of selective laser sintering

Mathieu Vandecasteele[1,*], **Ricardo Santander**[3], **Domenico Iuso**[2], **Mohsen Nourazar**[1], **Wilfried Philips**[1], **Sven Cornelissen**[3] and **Brian G Booth**[1]

[1] TELIN-IPI, Ghent University—imec, St-Pietersnieuwstraat 41, Gent B-9000, Belgium
[2] imec-Vision Lab, University of Antwerp, Universiteitsplein 1, Antwerpen B-2610, Belgium
[3] Materialise, Technologielaan 15, Leuven B-3001, Belgium
[*] Author to whom any correspondence should be addressed.

E-mail: mathieu.vandecasteele@ugent.be

Keywords: powder bed fusion, porosity, defect detection, selective laser sintering, *in-situ*

OPEN ACCESS

RECEIVED
15 August 2025

REVISED
11 December 2025

ACCEPTED FOR PUBLICATION
15 December 2025

PUBLISHED
30 December 2025

Original Content from this work may be used under the terms of the Creative Commons Attribution 4.0 licence.

Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

Abstract

Porosity formation in polymer selective laser sintering (SLS) significantly impacts part quality, as unwanted pores compromise mechanical properties. In-process control of porosity is a promising solution, but existing *in-situ* monitoring systems lack the necessary speed to meaningfully intervene in small localized areas, and the specificity to guide targeted process adjustments. This study addresses these limitations by introducing an *in-situ* monitoring system that predicts control-actionable porosity targets at high speeds (2 kHz). The system combines multimodal visible and short-wave infrared imaging, complementing information about the powder bed condition and thermal behavior. A neural network then makes predictions based on learned correlations between the image data and the porosity targets. The neural network incorporates both contextual and temporal information through a specialized architecture that leverages adaptive weight networks and a temporal sliding window. The proposed targets are extracted post-print using x-ray computed tomography analysis and serve as ground truth targets to train the neural network. Based on pore concentrations and spatial uniformity measurements, the targets have a clear physical justification, and allow any printed part to be used as training data. Crucially, by distinguishing between insufficient melting and overheating, these targets resolve the directionality of the required laser power adjustment, enabling non-ambiguous control actions. Experimental results on cylindrical test parts demonstrate good correlations between the porosity targets and expected pore formations. Furthermore, the proposed *in-situ* system shows strong predictive performance, with normalized mean absolute errors of 14% for predicting relative pore concentrations, and a correlation coefficient of 0.68 for predicting the underlying cause. The combination of a high-speed system and the prediction of control-actionable porosity targets pave the way for effective process control in SLS, improving part quality.

1. Introduction

Polymer selective laser sintering (SLS) excels in producing strong, nearly isotropic parts with good dimensional accuracy. It utilizes a laser to sinter powdered material layer by layer, fusing particles to form a solid structure. This process enables complex designs with excellent mechanical properties and durability, making it ideal for functional prototypes and end-use applications. SLS supports a wide variety of materials, offering flexibility for diverse applications. It achieves faster build rates than fused deposition modeling [1] and ensures efficient production of components.

However, as is common in all additive manufacturing (AM) technologies, polymer SLS is not devoid of imperfections in the final product, especially on the microstructural level [2]. The microstructure is primarily affected by the creation of pores (porosity) and unmelted residual particles, which prevent

© 2025 The Author(s). Published by IOP Publishing Ltd

full densification and cause anisotropic mechanical properties [3]. Even in optimal printing conditions, porosity levels of 3% to 8% are expected in polyamide 12 (PA12) printed parts [4, 5]. Given that PA12 and polyamide 11 (PA11) collectively dominate around 90% of the SLS materials market [6], effective porosity control is an important challenge.

Reducing porosity in SLS is inherently challenging because the production of consistent, high-quality parts requires a precise balance of laser parameters, such as laser power, scan speed, and hatch spacing, within a narrow optimal range to ensure complete melting and fusion [7–10]. If the energy input is too low, incomplete particle melting occurs, leading to porosity and poor layer adhesion [11]. Conversely, excessive energy can degrade the polymer, causing thermal decomposition and the formation of gas pockets within the material [11]. Porosity is also influenced by a complex interplay of additional factors, including but not limited to: powder particle morphology [4], build orientation [12, 13], part geometry [14, 15], heat accumulation [16], and powder age [17, 18]. The interplay of these factors makes it infeasible to define a universally optimal set of print parameters. As a result, the industry relies heavily on general guidelines [19] and expensive post-process quality inspections [20].

Recognizing the difficulty in defining universally optimal parameters, significant research efforts have focused on *a priori* porosity mitigation through systematic parameter studies and predictive simulations. Methodologies like design of experiments are frequently employed to map the complex relationships between printing parameters and resulting porosity levels [21, 22], helping to establish viable processing windows. However, due to the vast number of factors influencing the process, the task of understanding and optimizing all relevant factors is infeasible. Moreover, since not every factor can be controlled in these experiments, the conclusions are often narrowly applicable [2]. Parallel efforts involve computational modeling, including thermal simulations [23] and melt pool dynamics simulations [24], to predict temperature evolution and melt behavior, thereby identifying potential porosity hotspots during the design stage. While these simulation tools allow for virtual process optimization and design adjustments, their predictive power is inherently limited by model simplifications, computational cost, and, crucially, their inability to capture or react to stochastic variations within the powder bed or machine environment during an actual print [25].

A promising approach to mitigate porosity during SLS operation involves dynamically adjusting process parameters through in-process control. Enabling such control requires *in-situ* monitoring systems, which are becoming increasingly popular in AM [26–29]. These monitoring systems collect real-time data throughout the printing process, providing valuable insights into print quality. To drive an effective porosity-reducing control system, this data must identify targets that meet two criteria: (a) they must be estimated *in situ* using the real-time sensor data, and (b) they must provide actionable insights, not just indicating the presence of porosity but also pinpointing its likely causes to enable precise, corrective adjustments.

While the direct in-process control of porosity using actionable targets remains largely unexplored in published literature, prior work has targeted control of related process characteristics and *in-situ* defect detection. Notably, Phillips *et al* [30] developed a feed-forward automated laser control system utilizing infrared temperature measurements and a dynamic model to adjust laser power to improve temperature uniformity, reporting improvements of up to 57%. Building upon this, they later enhanced their approach [31] with intra-vector power modulation capabilities, further improving temperature uniformity. However, these approaches, while valuable for promoting global consistency, are inherently limited for precise, localized defect mitigation. Their primary focus on achieving overall thermal uniformity does not directly address the highly localized conditions that often lead to porosity formation, such as specific powder irregularities. Consequently, striving for global uniformity may fail to correct local issues causing pores [32].

Complementary to these control efforts, some research has focused on detecting porosity or density variations during or immediately after processing. For instance, Guan *et al* [33] employed optical coherence tomography to image both surface and subsurface features, distinguishing density variations of about 5% up to 300 μm depth at 180 frames per second (fps). In another approach, Schlicht *et al* [34] captured off-axis optical images at 1 fps. Image processing techniques were used to determine the normalized change of radiance of a part as a function of time, and deep learning was applied to estimate overall part porosity with a root mean square error of 0.6%. More recently, Hofmann *et al* [35] used off-axis thermal imaging (200 fps, 135 × 135 pixels) to extract handcrafted features and employed decision trees to predict voxel-level solidity, achieving a normalized mean absolute error of 3.36%. These methods show promise for *in-situ* quality inspection and could serve as digital twins to ensure consistency and reduce costly post-process checks. However, to truly integrate them into effective in-process control systems, they fall short in two major ways.

First, SLS operates at very high laser speeds (typically 0.25 to 5 ms^{-1}) meaning large areas are processed in a short time. Yet, local effects, such as geometry, non-homogeneous powder distribution [36], and non-uniform powder bed temperature [37], can induce porosity in very small regions. Meaningful control of porosity in these localized areas therefore requires rapid acquisition rates and fast signal processing algorithms, which current methods lack. Second, while current *in-situ* methods can predict an increase in porosity, they do not identify which underlying cause is responsible, making targeted interventions impossible. For example, excessive gas porosity might require lowering the laser power to minimize gas entrapment, whereas insufficient fusion necessitates increasing the power to ensure complete melting. Without the ability to correlate sensor input directly to control-actionable outputs (i.e. targeted process adjustments), these methods cannot recommend the precise interventions needed to address porosity formation. Hence, to effectively enable *in-situ* control of porosity, there is a need for high-speed monitoring systems that can correlate sensor inputs directly to control-actionable outputs in real time, enabling immediate adjustments during the printing process. This essential capability remains unaddressed in the current literature.

To address these shortcomings, this study proposes two main contributions. First, we define novel, control-actionable porosity targets derived from x-ray computed tomography (CT). Based on pore concentration and uniformity measurements, the novelty of these targets lies in their ability to distinguish between porosity caused by insufficient melting versus overheating. In contrast to standard scalar metrics which signal defect presence but leave the required adjustment open to interpretation, our proposed porosity targets enable non-ambiguous corrective actions, explicitly indicating whether energy input must be increased or decreased. Additionally, these targets can be extracted from any printed part, facilitating broader applicability for data-driven monitoring. Second, we develop a high-speed, multimodal *in-situ* monitoring system using visible (VIS) and short-wave infrared (SWIR) imaging at approximately 2 kHz. The combined data feeds into a GPU-optimized neural network designed for high-speed inference, integrating local contextual information such as part geometry to enhance accuracy.

In summary, the combination of both contributions helps to enable in-process control of porosity in the SLS process by:

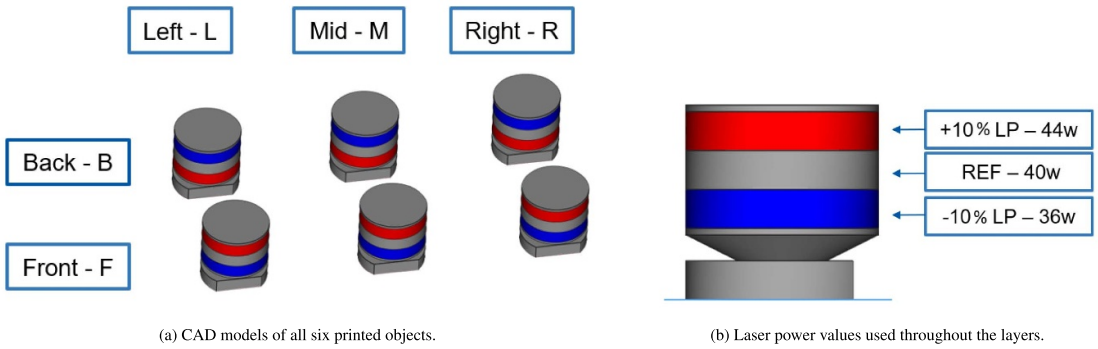
- Defining control-actionable porosity targets that serve as prediction outputs for *in-situ* monitoring systems, providing clear metrics for identifying and mitigating pore formation.
- Developing an *in-situ* monitoring system utilizing multimodal imaging and machine learning to accurately predict the proposed porosity targets at high speeds (> 2 kHz), enabling real-time adjustments.

To validate the proposed contributions, an experiment was performed with cylindrical samples using controlled intra-part laser power variations, recorded using the proposed *in-situ* system. Distinct porosity regimes were deliberately induced, ranging from insufficient melting to overheating, to generate a robust dataset for model training and evaluation. The remainder of this article is structured as follows. Section 2 details the experimental setup, describing the SLS system, the specific design of the test artifacts, and the hardware architecture of the *in-situ* monitoring system, including the justification for its design. Section 3 outlines the methodology, including the derivation of control-actionable porosity targets from CT analysis and the algorithmic design of the neural network pipeline. Section 4 presents the results, first validating the porosity targets and subsequently evaluating the predictive performance and inference speed of the monitoring system. Finally, section 5 discusses the implications of the findings, and section 6 concludes the study.

2. Experimental setup

2.1. SLS system

The experiment was conducted on an experimental SLS machine (Sindoh) with a usable build area of $450 \times 450 \times 500$ mm. The machine is equipped with two CO_2 lasers, each with a maximum power of 100 W and a wavelength of 10.6 µm. The material used was PA12. A two-stage preheating procedure preceded the printing process. The first stage involved heating the machine to 150 °C for one hour. During this stage, six layers of 'cold powder' at 120 °C were added at regular intervals. The second stage involved heating the machine to 173 °C for 30 min. In this phase, one layer of 'cold powder' was added every 20 s. The layer thickness was set to 100 µm. Following these preheating stages, the printing process commenced.



(a) CAD models of all six printed objects. (b) Laser power values used throughout the layers.

Figure 1. Visualization of the printed objects. (a) Computer-aided design (CAD) models of all objects with their labeling and color-indicated laser power. (b) Labeling of the laser power used for each zone. Red coloring indicates zones with 10% increased laser power (44 W), blue indicates 10% decrease in laser power (36 W), and grey indicates nominal laser power (40 W).

2.2. Experimental design

Six cylindrical samples were printed and analyzed, as illustrated in figure 1(a). All cylinders were printed in a single build, with their positions relative to the center of the build plate, as labeled in the figure. Each cylinder is divided into three zones, printed with distinct laser power settings: below-nominal, nominal, and above-nominal settings, as shown in figure 1(b).

The below-nominal zones are intended to induce porosity associated with insufficient energy input, while the above-nominal zones are designed to generate porosity due to excessive energy input. Specifically, a nominal laser power of 40 W (grey zones) was decreased by 10% for the below-nominal zones (blue) and increased by 10% for the above-nominal zones (red), as shown in figure 1(b)[4]. Nominal settings are defined as the optimal parameters derived from prior experience to minimize porosity for this specific geometry and powder.

Laser power was selected as the primary variable for this experiment due to its direct and significant influence on energy density and melt pool behavior, as well as the relative ease of precise control within a print. To isolate the impact of these controlled energy density deviations, other critical process parameters (scan speed, hatch distance, layer thickness, preheating conditions) were held constant at their established nominal values. While global machine settings are typically kept constant during printing, local thermal conditions within the part, influenced by factors like geometry or non-uniform build temperatures, can effectively alter the delivered energy density even when settings are unchanged. The controlled intra-part power variations in this study (+/− 10%) were designed to simulate these scenarios, specifically creating conditions intended to induce porosity associated with insufficient melting (below-nominal power) or potential overheating (above-nominal power). A deviation of 10% was selected based on preliminary experiments to balance two factors: it is sufficient to exceed the optimal process window, yet small enough to simulate realistic process fluctuations.

The nominal laser power of 40 W was selected to reflect established parameters routinely used in experienced SLS production environments for PA12, ensuring an industrially relevant baseline. This value is further validated by research on PA12 optimization [21], which reports ideal processing windows consistent with our settings. For the other laser parameters, the scanning speed was set to 4400 mm s^{-1} with a beam diameter of 0.6 mm, and a hatching distance of 0.3 mm. These parameters delivered an energy density of 0.0303 J mm^{-2} for the nominal zones. The variations in laser power resulted in energy densities of 0.0330 J mm^{-2} and 0.0273 J mm^{-2}, for the above-nominal and below-nominal zones, respectively.

For the cylinders in the front row (figure 1(a)), the sequence starts with the below-nominal zone, followed by the nominal zone, and ends with the above-nominal zone. The back row reverses this sequence to investigate the potential effects of thermal accumulation. For simplicity, cylinders will be referred to by their positional labels (e.g. cylinder 'MF' refers to the middle cylinder of the front row). Subsequent analyses focus on the top cylindrical section, beginning from the first non-nominal layer, thus excluding the base. Each cylinder has a total height of 15 mm, with the top cylindrical section having a diameter of 14 mm and a height of 10 mm.

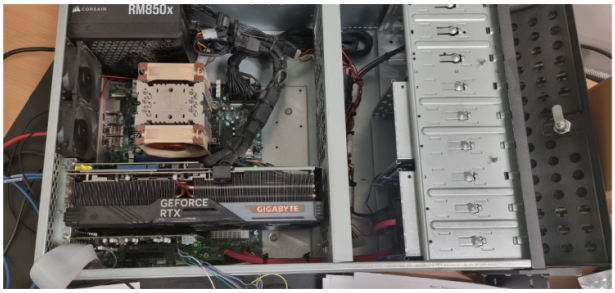
[4] Throughout this study, relevant figures will maintain the same coloring scheme to indicate these different zones.



(a) The experimental Sindoh machine.



(b) Positioning of the cameras in the Sindoh machine.



(c) The custom edge computing unit.

Figure 2. Pictures of the machine, hardware, and camera setup. (a) The experimental Sindoh machine. (b) The positioning of the cameras in the top module of the Sindoh machine. (c) The custom edge computing unit with integrated GPU.

2.3. *In-situ* system

The experimental print was completely recorded using our custom *in-situ* system. The system employs high-speed cameras in the visual and SWIR spectra wherein the processing is focused on the active region of melting (melt pool). The cameras are positioned off-axis, each tilted at approximately 20° relative to the build plate's normal. One camera, a prototype using a Luxima LUX2810 sensor, captured images in the VIS spectrum, while the other, a Xenics Cheetah 640 TE1, recorded in the SWIR range. The VIS camera operated at a frame rate of 2000 fps with a resolution of 96×96 pixels, with each pixel covering an area of around 150 µm $\times$ 150 µm. The SWIR camera operated at a frame rate of 1655 fps, with a resolution of 228×144 pixels, with each pixel covering an area of around 400 µm x 400 µm. Both cameras recorded at 8-bit depth.

To detect and control porosity in small local areas, the acquisition rate should be as high as possible to accommodate the fast-moving laser. However, as the acquisition rate increases, the maximum possible exposure time of the cameras decreases. With lower exposure times, less light enters the sensor and the signal-to-noise ratio decreases. This can be compensated by increasing external illumination, however, this is often hard to incorporate in existing SLS machines and might interfere with the powder bed heating lamps. In this work, it was found that 2000 fps provided a good trade-off between the achievable acquisition rate and the signal-to-noise ratio.

The physical integration is shown in figure 2, with the machine (figure 2(a)) and camera mounting (figure 2(b)) detailed specifically. The two cameras are strategically mounted above the processing area, viewing through dedicated windows. The VIS camera is mounted in front of a window featuring a low-pass filter designed to block the laser wavelength (around 10 µm). The SWIR camera is mounted in front of an unfiltered scanner window. The SLS machine is equipped with the Materialise control platform (MCP), which provides access to low-level control and machine data at 100 kHz, permitting the communication of laser positions and parameters. Data acquired from both the cameras and the MCP is processed by a custom edge computing unit, shown in figure 2(c). This dedicated unit, which includes an integrated NVIDIA RTX 4060 Ti GPU for high-speed computation, is responsible for synchronizing the camera frames with the laser position data and performing the algorithm's porosity inference computations. Since the cameras operate at different frame rates, image pairs were processed at the speed of the VIS camera, resulting in occasional duplicate SWIR frames within consecutive image pairs. For this study, the synchronized data was logged to disk for offline model training and validation, though

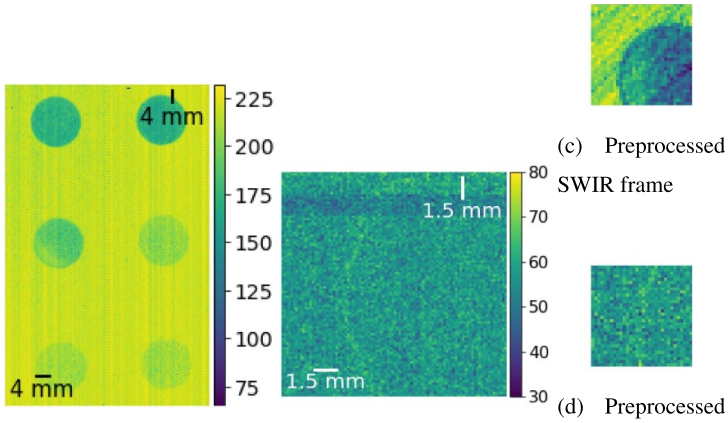


Figure 3. Example pseudo-colored unprocessed frames of (a) the SWIR camera and (b) the VIS camera; and processed frames of (c) the SWIR camera and (d) the VIS camera.

the platform is also designed with the capability to send control actions back to the MCP in a potential closed-loop configuration.

The system's design draws inspiration from our earlier work in selective laser melting (SLM) of metals [38, 39], which similarly employed high-speed VIS and SWIR cameras focused on the melt pool region. However, the sensing physics in SLS differs significantly from SLM. SLM operates at temperatures above 1000 K with vaporization temperatures in excess of 3000 K, causing intense thermal emission that results in a high spectral radiance in VIS/SWIR wavelengths [40], which clearly distinguishes the melt pool. In contrast, thermal radiance at typical SLS processing temperatures is negligible in these spectra. Therefore, the cameras in this SLS setup primarily capture interactions (reflectance, absorption) between external illumination and the powder bed or melt pool. Despite this difference, we hypothesize that these optical interactions are sufficient to reveal process phenomena correlated with porosity.

Figures 3(a) and (b) show pseudo-colored example frames during the printing of a layer of the cylinder, from the SWIR and VIS cameras, respectively. In the visual spectrum, the images mainly capture powder bed characteristics, such as surface roughness and packing density (compactness). This is visible in figure 3(b) which also shows the spherical boundary of one of the cylinders. These powder bed characteristics are known to change throughout a print and may induce porosity if non-optimal [41, 42]. Additionally, images in the visual spectrum have already been used in other work [34] to predict porosity phenomena with high accuracy. Thus, justifying the choice of sensor. The SWIR spectrum captures information more correlated to the temperatures of the powder bed and melt pool. It is well known that these temperatures must be very closely regulated during the SLS process, and any deviation from the optimal influences part porosity [43, 44]. The example frame (figure 3(a)) shows clear differences in intensity values for unsintered areas, sintered areas, and the powder bed. Unexpected increases or decreases in these values might indicate excessive or insufficient energy density, leading to porosity.

3. Methods

The method section is organized according to the two main contributions of this work. First, we define the porosity targets that will serve as prediction outputs of the *in-situ* monitoring system. These are derived post-process using x-ray CT analysis. Ultimately, the targets are needed as ground truth values for training the data-driven image processing algorithm to learn correlations between the *in-situ* sensor data and porosity formation. In contrast to existing works that primarily detect the presence of defects, the targets here are defined to be control-actionable, i.e. suggesting non-ambiguous corrective actions by resolving the directionality of the required process adjustment (e.g. increasing vs. decreasing energy density). Second, we present the algorithmic design of the *in-situ* monitoring system, describing the data processing pipeline used to predict these targets from high-speed sensor data.

3.1. Definition of porosity targets

To enable in-process control, porosity targets must characterize porosity in a control-actionable manner. Any control algorithm addressing porosity in SLS requires two types of metrics in real time:

1. How much porosity is the process creating compared to the nominal expected value?
2. Is the process printing with an energy density that is below-nominal, nominal, or above-nominal for this print job?

The first metric allows the control algorithm to determine if the process deviates from nominal conditions, thus requiring parameter adjustments. The magnitude of the metric indicates the extent of needed adjustment. Similar metrics were used in the earlier discussed *in-situ* systems of [33–35], but they lack the directional information necessary for parameter adjustment. This gap is addressed by the second metric, which indicates the current printing regime and unambiguously tells the control algorithm to either increase or decrease energy density if printing off-nominally.

To derive these metrics, the cylinders were analyzed using high-resolution x-ray CT. All cylinders were scanned with the UAntwerp FleXCT scanner [45]. Scans were performed at 120 kV, with 1.5 mm Aluminum filtration, 1266 projections with 174.02 mAs exposure, a 250 mm source-to-object distance, and a 750 mm source-to-detector distance. The scanning process yielded a resolution of 10 µm in each of the three dimensions. The CT slices were aligned with the CAD models such that the Z-axis corresponds to the build direction of the cylinders. A segmentation technique similar to [46] was used to segment and characterize individual pores in 3D from the CT data. Based on this segmented data, the two porosity targets are defined below.

We base our first porosity target on pore concentration, which is defined as the number of pores per cubic millimeter. This metric was chosen based on studies [12, 47] that demonstrate CT pore concentrations are more sensitive to off-nominal printing regimes compared to other commonly used measurements, such as porosity percentage. Additionally, these studies report that pore concentration exhibits a stronger correlation with the mechanical properties of printed parts.

The pore concentration C_p is defined as the number of pores $N\mathrm{p}$ within a volume V divided by that volume: $C\mathrm{p} = N\mathrm{p}/V$. A pore belongs to the volume if its center of mass lies within it. Since some porosity is always expected, even in nominal zones, the measurement is expressed relative to some expected concentration $\mathbb{E}[C\mathrm{p} \mid nom]$. This denotes the expected concentration of pores under nominal printing conditions. The final porosity target is referred to as the relative pore concentration $\widetilde{C}\mathrm{p}$ given as:

$$\widetilde{C}\mathrm{p} = \mathrm{ReLU}\left(\frac{C\mathrm{p} - \mathbb{E}\left[C\mathrm{p} \mid \mathrm{nom}\right]}{\mathbb{E}\left[C\mathrm{p} \mid \mathrm{nom}\right]}\right) \tag{1}$$

where ReLU denotes the rectified linear unit function: $\mathrm{ReLU}(x) = \max(0, x)$. For ease of interpretation, the target is zero or close to zero if the pore concentration is as expected in a nominal printing regime, and the value increases in off-nominal printing regimes. If the concentration is less than the nominal expected value, then no control action is needed; thus, the output remains zero.

The second porosity target is derived from spatial distributions of pores within single print layers. It is based on the observation that partially or non-molten particles have a tendency to align within a layer [12, 48]. The underlying cause is believed to be a temperature gradient along the depth of the melt pool: the laser heats the upper particles more than the lower particles, resulting in lower temperatures and potentially incomplete melting at the bottom of the melt pool [49, 50]. Pavan *et al* [51] observed that this temperature gradient can induce periodic variations in part density, which can be detected across CT slices. This density periodicity is also evident in our experimental data, as seen on figure 4, for the below-nominal zone of the MF cylinder. This figure shows the average intensity (a proxy for density) of individual CT slices, where the average was computed from solid-material voxels (excluding pores). The effect can also be observed on figure 5, that shows the CT scan of a cross-section of the below-nominal zone of the MF cylinder. It clearly shows horizontal streaking patterns that correspond to the CT intensity periodicity as calculated in figure 4.

We hypothesize that in below-nominal printing zones, insufficient melting at the bottom of the print layer not only reduces local density but also increases the likelihood of pore formation. Consequently, it is expected that pore concentration distributions in below-nominal printing zones are not uniform across a printing layer. Instead, pores are more concentrated in areas corresponding to lower density zones. If true, and the same pattern is not observed for nominal or above-nominal printing zones, quantifying this relationship would provide a metric that correlates porosity patterns to specific printing regimes.

We quantify this relationship using the Pearson correlation coefficient. In particular, the correlation coefficient is computed between the distribution of the pore concentration within a single print layer L: $C\mathrm{p}(L)$, and the density distribution of that print layer $S(L)$. A print layer L is represented as an ordered set of CT slices $l_i : L = (l_1, l_2, \ldots, l_n)$ where l_1 represents the bottom CT slice and l_n the top slice. The

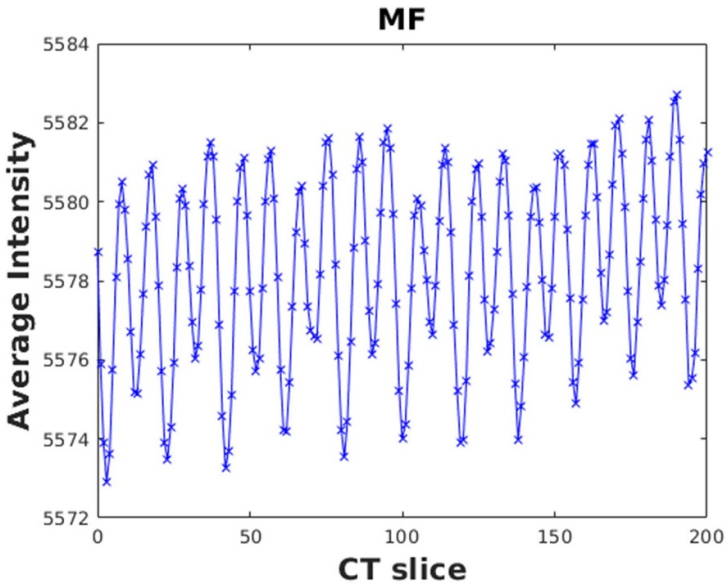


Figure 4. Average CT intensity as a function of the CT slice index for the MF below-nominal zone.

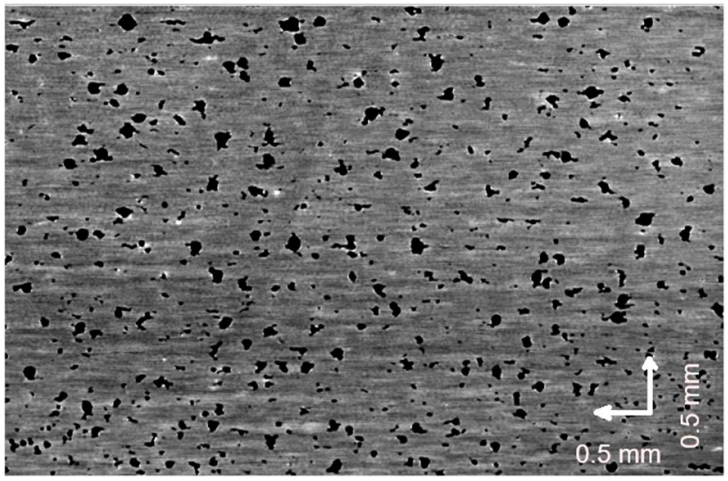


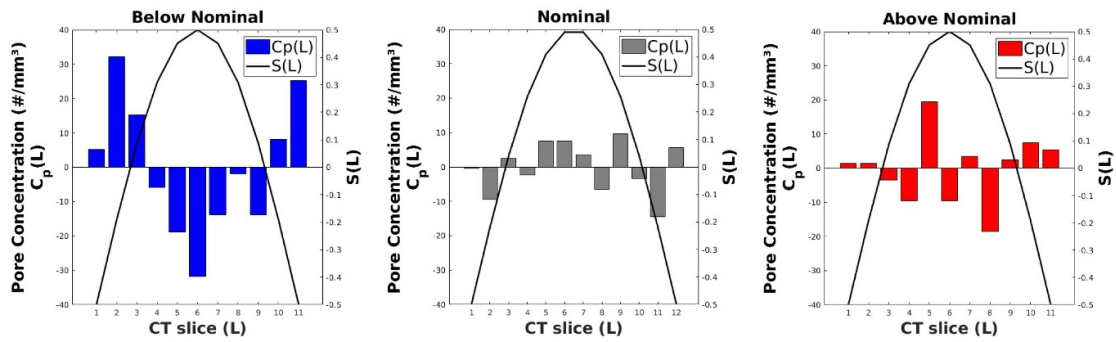
Figure 5. CT cross-section of the below-nominal zone of the MF cylinder. The horizontal streaking indicates periodic density variations attributed to the laser-induced temperature gradient.

second porosity target is then calculated as follows:

$$\rho_{Cp(L),S(L)} = \frac{\mathrm{cov}\left(Cp\left(L\right), S\left(L\right)\right)}{\sigma_{Cp(L)}\ \sigma_{S(L)}} \tag{2}$$

Where cov is the covariance function and σ denotes standard deviation. Since $\rho_{Cp(L),S(L)}$ roughly relates to how uniform the pore distribution is within a single layer, it will further be referred to as the pore uniformity target.

Figure 6 illustrates the relation between $Cp(L)$ and $S(L)$ for some randomly chosen layers of the MF cylinder for each zone. Note that the mean is subtracted from both quantities, an operation part of the covariance calculation. In the below-nominal layer (figure 6(a)) the sign of $Cp(L)$ is mostly the opposite of $S(L)$ for each l_i and the amplitudes display the same pattern. This indicates a strong negative correlation, which for this particular example produced a correlation coefficient of −0.81. Such a pattern is absent in other zones (figures 6(b) and (c)), where the correlation is weak: −0.15 for the above nominal zone, or even exhibits more positive correlations: 0.43 for the nominal zone. In general, it is expected that for the below-nominal layers, the coefficient is closer to –1, and closer to 1 for nominal and above-nominal layers.



(a) Below-nominal zone. (b) Nominal zone. (c) Above-nominal zone.

Figure 6. Example illustrations of the Pearson correlation coefficient calculation for three randomly chosen layers of the MF cylinder from the (a) below-nominal, (b) nominal, and (c) above-nominal zones.

Algorithm 1. Printing Layer Identification.

Input: l = Ordered set of all CT slices
$I(l)$ = Average intensity of each CT slice l_i
P = Nominal expected number of CT slices per print layer
Output: K = Index subset of l where each element indexes the beginning of a print layer

$K \leftarrow$ FindLocalMinima $(I(l))$
$K_{temp} \leftarrow K[0]$
for i = 0..len(K) **do**
  $D \leftarrow$ PairwiseDistance $(K_{temp}[-1], K[i:-1])$
  $N =$ argmin $(|D - P|)$
  $K_{temp} \leftarrow K_{temp} \cup [N]$
end for
$K_{temp} \leftarrow$ Filter $(K_{temp} < P/2)$
$K \leftarrow$ Insert $(K_{temp} > 1.5P)$

Since both $Cp(L)$ and $S(L)$ are distributions over CT slices within single printing layers L, calculating the pore uniformity target requires identifying the physical printing layers L within the CT volume. Even though the nominal print layer thickness is fixed (100 µm) throughout the print, the effective layer thickness of each layer can deviate due to changing powder properties and build chamber conditions [52]. Consequently, assuming a fixed layer thickness to identify print layers in the CT scan would be inaccurate.

Instead, print layers can be identified by analyzing the density patterns [51] (e.g. figure 4). It is assumed that a valley-to-valley period corresponds to a single print layer. Figure 4 suggests that detecting local minima suffices for layer identification. However, due to noise in the density signal for some cylinders and zones, a robust minima detection method is required. In this study, print layers were identified in the CT volume using the method described in the abstracted pseudo-code in Algorithm 1, illustrated as a flowchart in figure 7.

The *FindLocalMinima* function identifies the CT slices l_i that satisfy the condition $I(l_i) < I(l_{i-1}) \wedge I(l_i) < I(l_{i+1})$. The algorithm then proceeds with the very first local minimum and does a greedy search for the next local minimum, selecting the location closest to the nominal period P (print layer thickness / CT scanner resolution). After the greedy search, local minima are removed if they lie closer than half the nominal period to another minimum. Finally, gaps larger than $1.5P$ are removed by inserting points in the middle. All CT slices $K_n \leqslant l_i < K_{n+1}$ are then considered to be CT slices part of a single printing layer L_n. Since the density pattern $S(L)$ can be noisy, it is instead modeled by half a period of a sinusoid, wherein the sinusoid has a period of $2|L|$ and is sampled at CT slices l_i: $S(L) = \sin\left((\pi\ /\ |L|)\, l_i\right)$.

The porosity targets, relative pore concentration and pore uniformity, can be extracted per print layer from any arbitrary part using post-process CT analysis, serving as ground truth for *in-situ* monitoring. Importantly, the targets are control-actionable: the relative pore concentration indicates the extent of porosity being created, allowing for the identification of off-nominal printing. Meanwhile, pore uniformity identifies the printing regime responsible, enabling non-ambiguous control actions (e.g. lowering or increasing energy density).

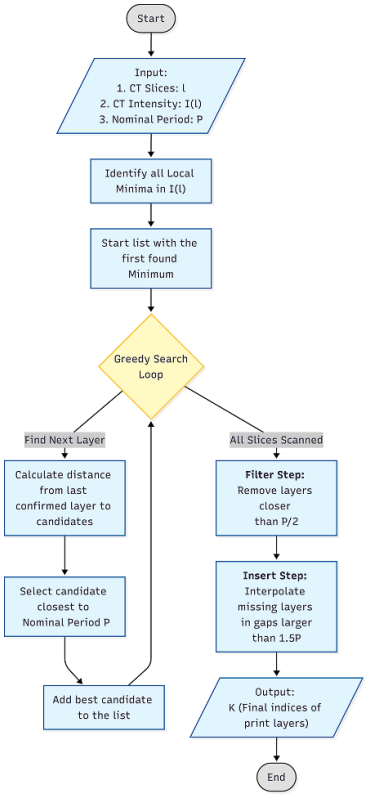


Figure 7. Flowchart representation of algorithm 1, describing the procedure to find individual print layers in the CT scan.

3.2. Algorithmic design

The algorithmic framework of the *in-situ* monitoring system is detailed in the following four subsections. First, *Data Preprocessing* covers the processing steps performed on the captured images to enhance throughput and reduce unnecessary variance. Second, *Feature Extraction* describes the projection of these images into a compact lower-dimensional feature space using a variational autoencoder (VAE). Third, *Porosity Prediction* details how these features are combined with local process context and temporal history to predict the porosity targets. Finally, *Data and Training Procedure* outlines the dataset composition, integration of ground truth labels, and training protocols employed to optimize the neural networks.

3.2.1. Data preprocessing

Prior to further processing, the frames from both cameras are cropped into smaller 32×32 pixel areas, centered on the position of the laser, to accelerate processing. The mapping from laser coordinates to VIS pixel coordinates, necessary to find center positions to crop around, was calculated by a 4-point calibration based on the printing and recording of a simple square calibration pattern. The mapping from laser coordinates to SWIR image coordinates was achieved by finding an optimal transformation using the iterative closest point algorithm [53] between cylinder pixels segmented from one of the images and the corresponding laser coordinates. Additionally, the frames are rotated to always have the same orientation relative to the laser direction, to avoid irrelevant variance induced by it. Figures 3(c) and (d) show the preprocessed versions of the SWIR and VIS frames, respectively.

3.2.2. Feature Extraction using VAEs

The preprocessed images are passed through feature extractors, transforming the images into lower-dimensional representations. These feature extractors are trained separately from the rest of the model, using fully-connected VAEs. Herein, a probabilistic encoder, $e_\phi(\mathbf{x})$, compresses the input image $\mathbf{x}$ into $\mathbf{z}$ (feature space), where it is modeled by a multivariate Gaussian probability distribution parameterized by a mean vector $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$: $\mathbf{z} = \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$. The decoder $d_\theta(\mathbf{z})$ then maps the feature space back to the output image $\hat{\mathbf{x}}$. A separate VAE is trained for each modality on a subset of the images.

The VAE architecture and training procedure are adapted from [54]. The structure of the model is illustrated in figure 8. The VAEs are trained using the following objective function (equation (3)):

$$L(\mathbf{x}, \hat{\mathbf{x}}) = ||\mathbf{x} - d_\theta(e_\phi(\mathbf{x}))||_2^2 + \beta \, \mathrm{KL}(e_\phi(\mathbf{z}|\mathbf{x}) \, || \mathcal{N}(\mathbf{0}, \mathbf{I})). \tag{3}$$

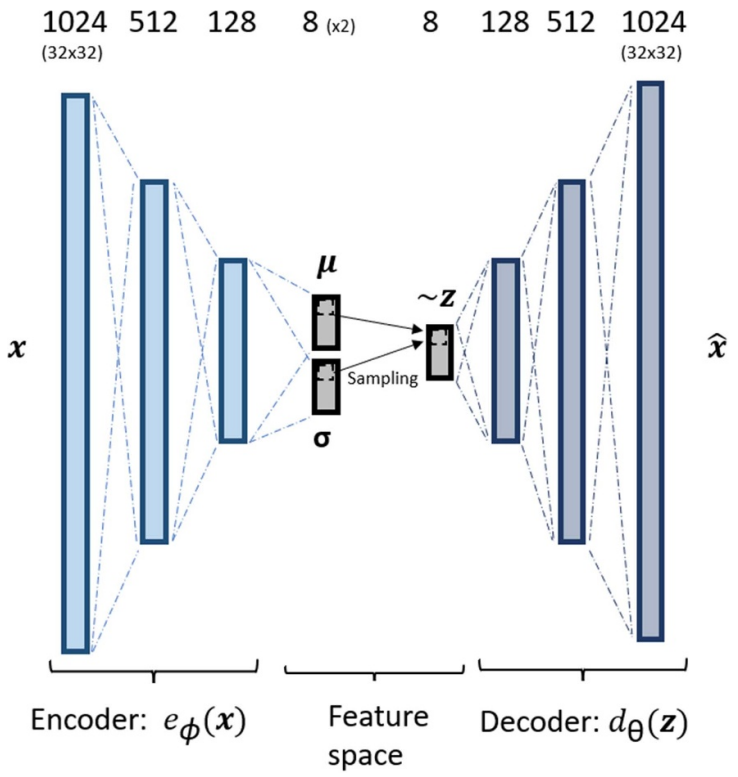


Figure 8. Architecture of the VAEs.

As opposed to standard VAEs, we employ β-VAE [55] combined with the ControlVAE technique [56], utilizing a non-linear proportional-integral (PI) controller to steer the Kullback–Leibler (KL) divergence loss towards a predefined target value. KL( ) denotes the KL-divergence distance, and β the scalar controlled by the PI-controller to steer this distance towards the target value. During inference, only the encoder $e_\phi(\mathbf{x})$ of the VAE is utilized, with the mean vector $\boldsymbol{\mu}$ serving as the final feature vector for each image. Each encoder produces 8 output features per image modality, $\boldsymbol{\mu} \in \mathbb{R}^8$.

3.2.3. Porosity prediction using context-adaptive neural networks

The overall architecture (figure 9) is inspired by our earlier work [39, 57] and comprises three main components: the VAE-based feature extractors, an adaptive weight network (AWN), and a porosity prediction network.

The AWN [58] combines the image features with contextual information from the printing process termed auxiliary inputs, which capture dynamic geometric and process-related variables. Let $\mathbf{a}_t$ denote a vector of these auxiliary inputs at timestep t. The AWN then employs a neural network with parameters $\mathbf{W}_g$ to implement a function $G(\mathbf{a}_t;\mathbf{W}_g) : \mathbb{R}^d \rightarrow \mathbb{R}^{m\times n}$, which outputs a matrix of scalars. This matrix is combined with the concatenated image features (output of the VAE encoders), denoted by $\mathbf{f}_t$, via vector-matrix multiplication. This process yields contextually-modified features, $\mathbf{f'}_t$: $\mathbf{f'}_t = G(\mathbf{a}_t;\mathbf{W}_g)\,\mathbf{f}_t$. Functionally, the output of the AWN acts as a fully connected layer (without bias or activation), where the input is the image features $\mathbf{f}_t$, and the output is the contextually adapted features $\mathbf{f'}_t$. However, unlike a traditional fully connected layer, the weights are dynamically adjusted based on the auxiliary inputs $\mathbf{a}_t$.

In addition to [39], temporal information was incorporated. This involves maintaining a buffer of previous encoded frames and fusing the contextual information with the entire buffer, instead of a single encoded VIS-SWIR image pair. Let B denote an ordered list that contains SWIR and VIS concatenated features $\mathbf{f}_i$ of 3 previous timesteps and the current timestep: $B = (\mathbf{f}_t, \mathbf{f}_{t-1}, \mathbf{f}_{t-2}, \mathbf{f}_{t-3})$. Let $\mathbf{f}_{t:t-3}$ denote the concatenation of the elements of B: $\mathbf{f}_{t:t-3} = \text{concat}(\mathbf{f}_t, \mathbf{f}_{t-1}, \mathbf{f}_{t-2}, \mathbf{f}_{t-3}) \in \mathbb{R}^{64}$. The final (redefined) context-adapted features $\mathbf{f'}_t$ are then computed as:

$$\mathbf{f'}_t = G\left(\mathbf{a}_t;\mathbf{W}_g\right)\ \mathbf{f}_{t:t-3}. \tag{4}$$

The buffer is easily adjusted at each time step by shifting it 16 positions and inserting the new $\mathbf{f}_t$. This approach can be thought of as a temporal sliding window method.

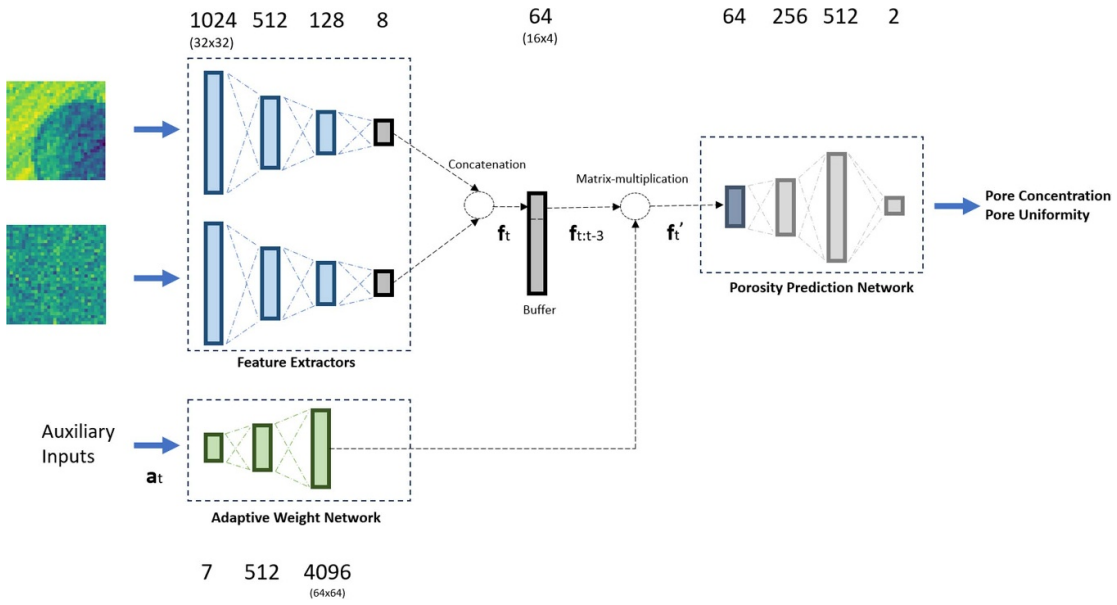


Figure 9. Architecture of entire porosity prediction framework.

Table 1. Used auxiliary features and the corresponding physical printing properties they capture.

Description	Captures
Minimum within-layer distance between the laser position and the object border	Border and corner effects
Minimum within-layer distance between the laser position and the closest overhanging region	Closeness to overhang geometries (horizontally)
Minimum within-layer distance between the laser position and the closest bulk region	Closeness to bulk region if within overhang region
Length of the current scan line	Thickness of part, thermal gradients
The direction of laser travel as a $(\Delta X, \Delta Y)$ tuple	Laser direction
Hatching pattern orientation	Hatching orientation

The central idea is that the optimality of the process is highly dependent on the local printing context [59]. Including local contextual information, using AWNs, is both essential for accurate defect prediction and more efficient than relying on the model to implicitly learn these complex dependencies. Table 1 summarizes the contextual inputs employed, and the physical printing properties they capture. The inputs are defined in more detail in [60]. These auxiliary inputs are precomputed and stored in look-up tables, enabling rapid retrieval for high-speed inference.

The porosity prediction network maps the contextually modified feature vector $\mathbf{f'}_t$ to the porosity targets. Throughout the architecture, layers are fully-connected, followed by batch normalization and ReLU activation, except for the final layers of the feature extractors, the AWN, and the porosity prediction model. The conditioned feature vector $\mathbf{f'}_t$ is kept at the same size as $\mathbf{f}_{t:t-3}$, therefore $G(\mathbf{a}_t; \mathbf{W}_g)$ must output a 64×64 matrix. To achieve this, the AWN's final layer contains 4096 units, which are reshaped into a 64×64 matrix before multiplication with $\mathbf{f}_{t:t-3}$.

3.2.4. Data and training procedure

The complete experimental build (6 cylinders) was recorded, yielding a total of 243 682 synchronized VIS-SWIR image pairs across 132 layers during active laser exposure. For each VIS-SWIR image pair, the associated auxiliary features were computed (see table 1) based on the laser's position. These image pairs, combined with the auxiliary features, constitute the network input. Each input sample is paired with corresponding CT-derived ground truth values. Specifically, a tuple containing the relative pore concentration and pore uniformity. The input-output pairs were utilized to train the neural network (figure 9) and subsequently validate its predictive performance.

To ensure a representative assessment, separate training and test sets were created from the total dataset. From the total set of cylindrical layers (87 layers), every third layer in the sequence was assigned

to the test set, and the remaining layers to the training set. In total, the training set consisted of 118 192 samples, while the test set consisted of 59 446 samples.

Model training involved two separate steps. First, the VAEs (figure 8) were trained independently to function as feature extractors for the VIS-SWIR image pairs. Second, the AWN and prediction model layers were jointly trained, using the extracted VAE features and the auxiliary features to predict the porosity targets.

The VAEs were trained using 18 randomly selected layers, totaling 32 247 image pairs. Of those 18 layers, 6 layers were chosen from the bottom non-cylindrical base part of the cylinders to increase variance in the image space and produce more robust image features. The VAEs were optimized by minimizing the loss function defined in equation (3) using the Adam [61] optimizer. The learning rate was set to 0.0001, which decreased with an exponential rate of $\gamma = 0.99$ for each epoch, using a batch size of 64 for 100 epochs. The KL-divergence target loss was set to 30.

The AWN and prediction network layers were trained jointly, utilizing both auxiliary features and image feature buffers (VAE outputs) to predict the porosity targets. The training employed a mean squared error loss function and the Adam optimizer. A batch size of 64 was used, with a learning rate initialized at 0.00001 and an exponential decay rate of $\gamma = 0.99$ applied per epoch, over a total of 100 epochs. This training phase utilized 118 192 samples.

4. Results

This section first assesses whether the proposed porosity targets capture the intended off-nominal conditions introduced in the print experiment. Subsequently, we evaluate the *in-situ* monitoring system and trained neural network in terms of prediction performance on the porosity targets of the test set. An ablation study is also performed to assess the value of the AWN and temporal processing. At last, we test the inference time of the entire system, to verify the real-time requirements.

4.1. Porosity targets

Figure 10 shows pore concentrations for each cylinder and each zone. Increases (relative to nominal zones) are observed in the below-nominal zones of the back row, with an average increase of 112.9%; and the below and above-nominal zones of the front row with increases of 61.7% and 61.0% respectively. Pore concentration was decreased by 5.3% for the above-nominal zones of the back row.

A key observation from figure 10 is the contrasting effect of the +10% power deviation: while it increased porosity in the front row cylinders (LF, MF, RF) by 61.0% relative to nominal, it unexpectedly decreased porosity by 5.3% in the back row cylinders (LB, MB, RB). This difference strongly suggests an interaction between the applied laser power and the evolving thermal conditions within the parts, driven by the intentionally reversed printing sequence between the rows. Additionally, while the physical separation between the rows is small, slight variations in laser incidence angle between these locations may have contributed to minor differences in the delivered energy density [62], further compounding these thermal effects.

Specifically, for the back row cylinders, the above-nominal zones were printed first onto a relatively cooler base with little accumulated heat. In this state, the effective energy density likely brought the process conditions closer to the optimum for that specific cooler thermal state, reducing porosity compared to nominal. Conversely, for the front row cylinders, the above-nominal zones were printed last onto a significantly hotter base due to accumulated heat from preceding zones. Here, the same laser power resulted in a higher effective energy input, pushing the process into an overheating regime and thus increasing pore concentration. This observation is consistent with how nominal process parameters are typically established: often, a parameter set is chosen that optimizes the average porosity over an entire part. This method implicitly accounts for thermal accumulation, as later, hotter layers influence the outcome. Consequently, the chosen nominal setting, while optimal on average, might be slightly suboptimal for the initial, cooler layers (consistent with the +10% power improving results in the back row) and potentially closer to the overheating threshold for the final, hotter layers.

The calculation of the relative pore concentrations (equation (1)) for each printing layer and part is shown in figure 11. The nominal expected pore concentration $\mathbb{E}[Cp \mid nom]$ was calculated by averaging the pore concentration values for all nominal layers. Overall, the expected pattern is observed, with the largest values seen in zones previously identified with elevated pore concentrations (figure 10). Some small values appear for nominal zones but are expected due to natural variance. Exceptions are notable for the front-row cylinders just before the onset of the above-nominal zone. Notably, a ramp-up in relative pore concentration precedes the above-nominal zone. It may be the case that the printing in the

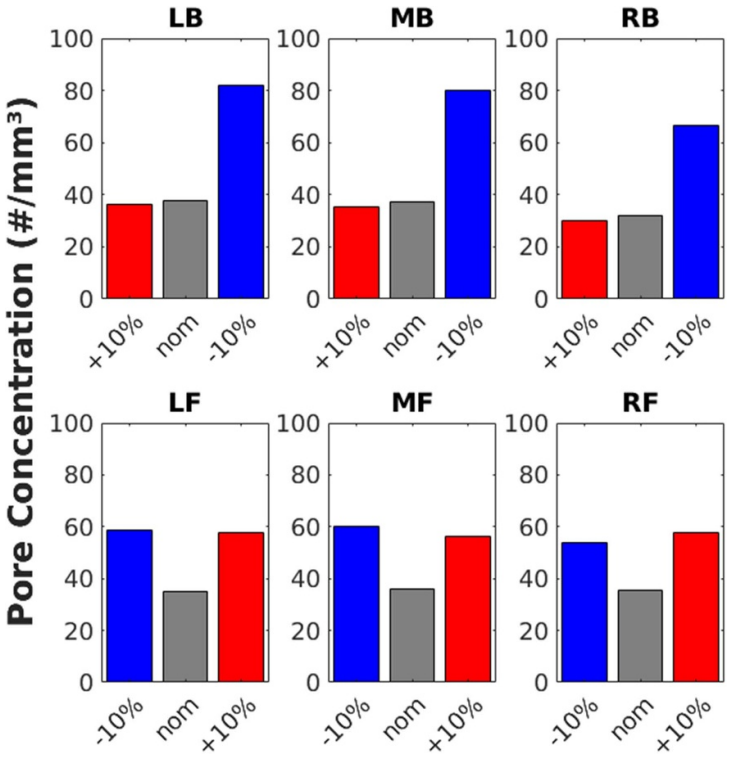


Figure 10. Pore concentration for each zone of each cylinder.

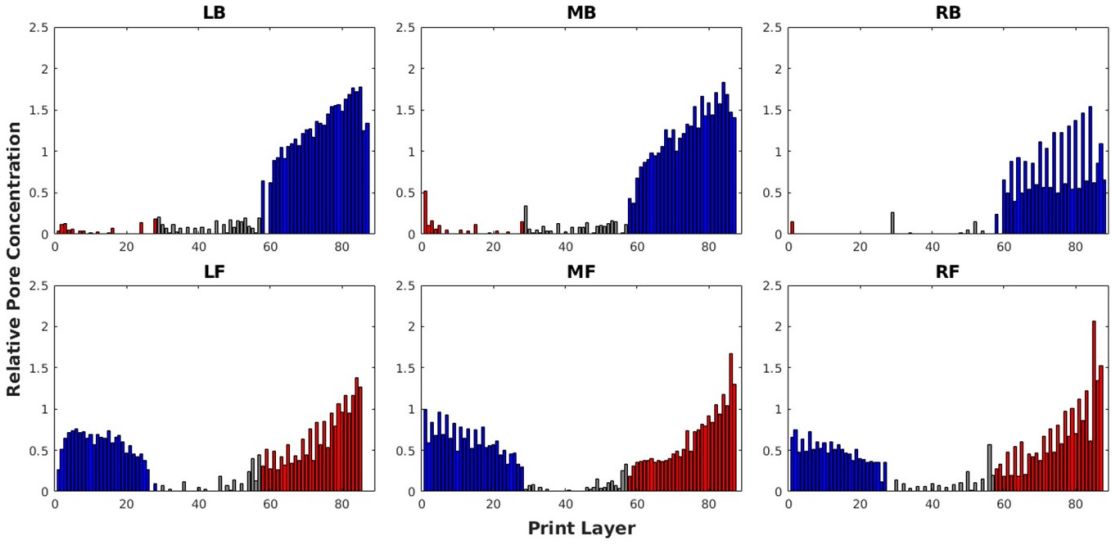


Figure 11. The final relative pore concentration values as a function print layers. Red coloring indicates layers with 10% increased laser power, blue indicates 10% decrease in laser power, and grey indicates nominal laser power.

nominal zone caused thermal build-up, leading to excessive temperatures near the end of the zone. The upper off-nominal zones show an increasing trend, indicating that the printing conditions are worsening during the print. Meanwhile, the below-nominal zones in the front row shows a somewhat opposite trend, with the relative pore concentrations decreasing with layer height. This too could be an indication of thermal buildup, but this trend is not as clear as the others.

Figure 12 shows the pore concentration, $Cp(L)$, per CT slice overlaid with the slice-averaged CT intensities for the different zones of the MF cylinder. In the calculation of $Cp(L)$, pores with a volume smaller than 50 voxels were excluded, to isolate pores more likely to be caused by deviating laser parameters. For the below-nominal zone (figure 12(a)), the increases in pore concentration align with the valleys of the slice-averaged CT intensity signal. As hypothesized, this leads to the conclusion that for

below-nominal printed zones, the pores are mainly concentrated within the lower-density regions. The same cannot be observed, or at least not to the same extent, for the nominal and above-nominal zones (figures 12(b) and (c)). In these zones, the pore concentration is more uniform and does not show any clear correlations with the slice-averaged CT intensities. The same observations were made on all cylinders.

The values for the pore uniformity (equation (2)) are visualized in figure 13. The figure shows the unaltered values (blue), as well as a moving average with a window of size 5 (red). It is evident that the unaltered values are very noisy. This is not surprising since porosity patterns in SLS have been shown to be influenced by many factors that produce multiple compounding periodical effects [21, 51]. Nevertheless, the moving average confirms our hypothesis: the below-nominal zones show values that are closer to –1, and the other zones generally have values closer to 0.

Due to the noisiness of the unaltered values, the averaged values were used instead for the final ground truth porosity targets. To simplify the interpretation, an offset value of −0.2 was used such that values above 0 are associated with nominal or above-nominal zones, and values below 0 are associated with below-nominal zones. The final used targets are shown in figure 14. These targets largely align with expectations based on the experimental design. Some notable exceptions are the below-nominal zone of the LF cylinder and some deviations at the transitions between zones.

4.2. Performance of the *in-situ* monitoring system

4.2.1. Predictive performance of the model

The entire model (figure 9) was optimized on the training set, mapping each VIS-SWIR image pair and its corresponding auxiliary features to a tuple of relative pore concentration and pore uniformity (derived from figures 11 and 14). The model was then evaluated on the test set.

Figures 15 and 16 show predicted values for every test layer of each cylinder, for the relative pore concentration and pore uniformity respectively. The predicted values for all samples within a single layer were aggregated by averaging. It is evident from figure 15 that the relative pore concentration predictions follow the same pattern as the ground truth values. For the back row cylinders, the predictions start low and then experience a sharp increase at the onset of the increased porosity due to below-nominal printing. The front row predictions begin with increased values due to the below-nominal printing zone, subsequently drop in the middle nominal zone, and finally increase again with the start of the above-nominal zone.

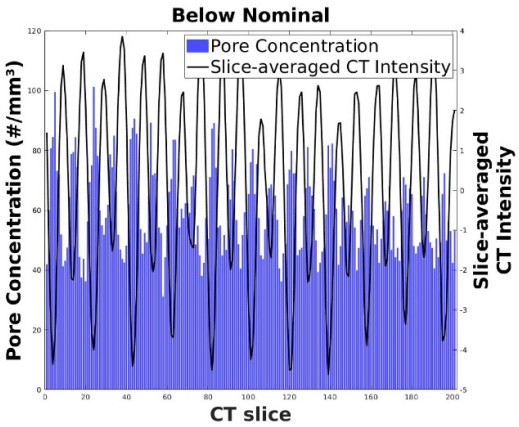
The predictions for the pore uniformity (figure 16) also follow the expected pattern. Values for the back row cylinder generally remain above 0, until the start of the below-nominal printing zone, at which point the averaged predictions drop below 0. The opposite pattern is observed for the front row cylinders: values start below 0 (corresponding to the below-nominal zone) and subsequently increase above 0 at the onset of the nominal zone.

Quantitative performance metrics are listed in tables 2 and 3 for the relative pore concentration and pore uniformity predictions respectively. For the relative pore concentration, the coefficient of determination R^2 was computed between the averaged layer predictions and the ground truth targets. For the pore uniformity, the Pearson correlation coefficient ρ was computed instead, since for those results, it is more important that the predictions have the same sign, and the amplitudes are less of a concern. For both targets, the normalized mean absolute error (NMAE) was computed using equation (5):

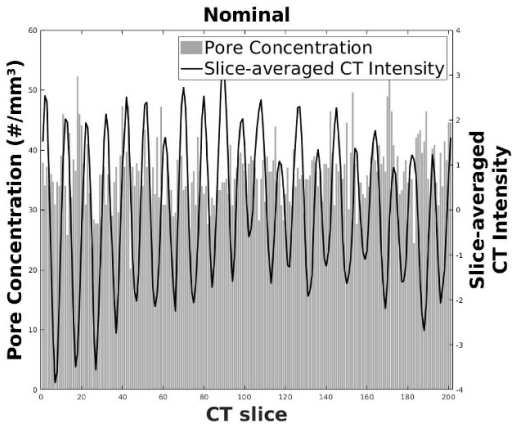
$$\mathrm{NMAE} = \frac{\sum_i^n |x_i - y_i| / n}{\max(\mathbf{y}) - \min(\mathbf{y})} \tag{5}$$

where the mean absolute error between predicted values x_i and ground truth values y_i is normalized by the range of the entire set of ground truth values $\mathbf{y}$. The tables show the evaluation metrics for each cylinder separately as well as the average over all cylinders.

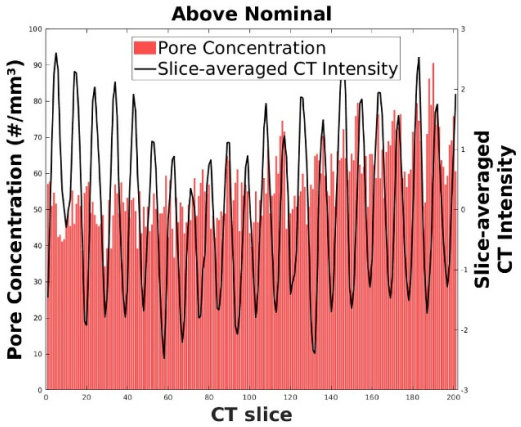
The highest R^2 values are observed (table 2) for the back-row cylinders with values above 0.7. These values drop somewhat for the front-row cylinders, especially the LF cylinder. These differences are also noticeable in figure 15, where the transition from off-nominal to nominal zones is not as significant for the front-row cylinders as in the back-row cylinders. The lowest value for the LF cylinder is attributed to the noisy predictions of the nominal zone and the large difference in magnitude for the below-nominal zone. The R^2 metric captures the overall fit between the averaged predictions and ground truth values. However, the prediction error at the level of individual data points is perhaps more critical for assessing model performance. This is quantified by the NMAE metric, which averages 0.14 across all cylinders, indicating moderate to strong predictive capabilities [63].



(a) Below-nominal zone.



(b) Nominal zone.



(c) Above-nominal zone.

Figure 12. Relationship between the pore concentration values and slice-averaged CT intensity values for each CT slice of the MF cylinder for (a) the below-nominal zone, (b) the nominal zone and (c) the above-nominal zone.

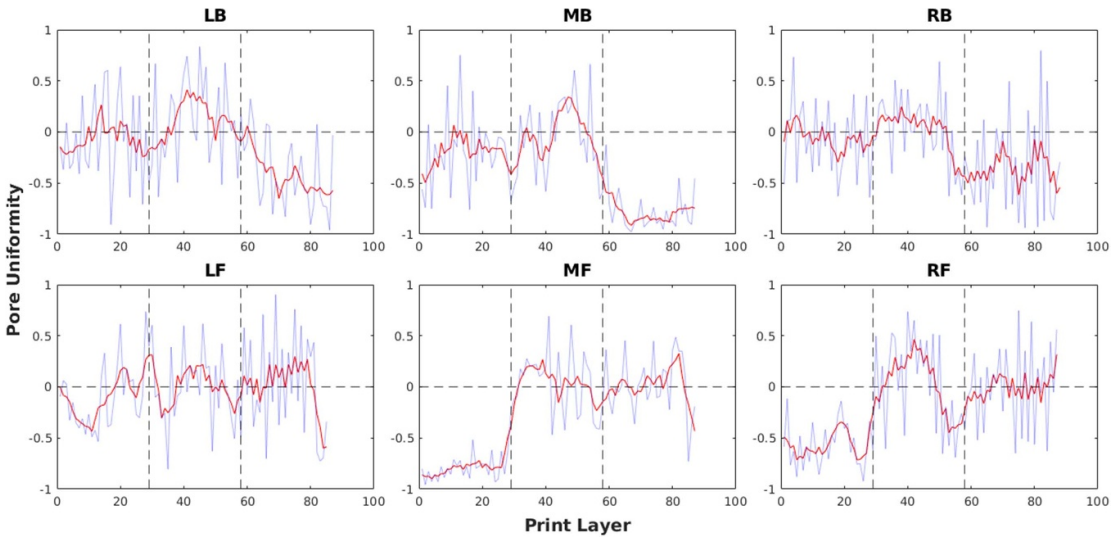


Figure 13. Pore uniformity as a function of print layer. Blue traces indicate unaltered values, while red traces show a moving average (window size 5).

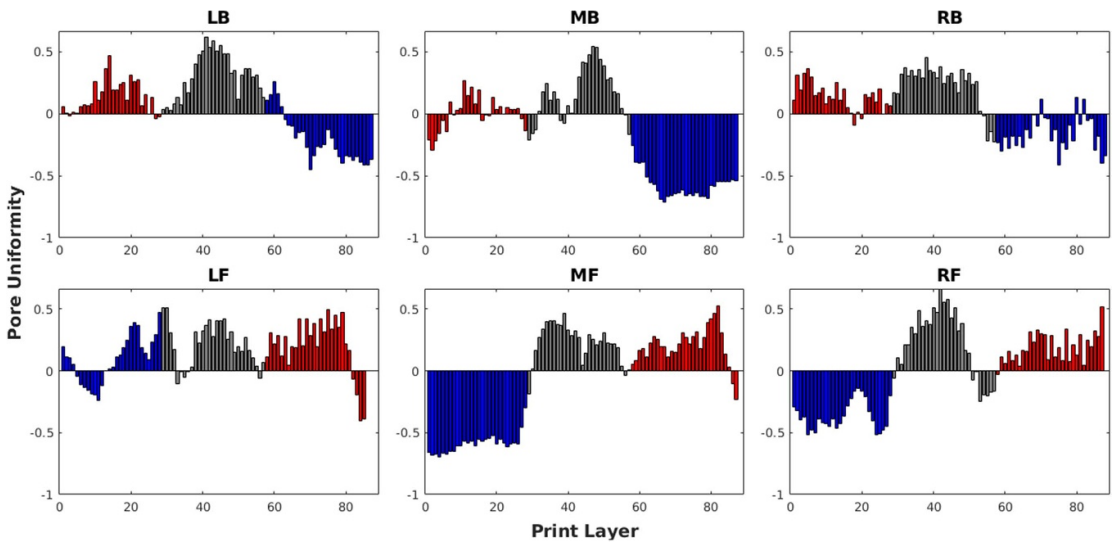


Figure 14. Final averaged and offset pore uniformity targets as a function of the print layers. Red coloring indicates layers with 10% increased laser power, blue indicates 10% decrease in laser power, and grey indicates nominal laser power.

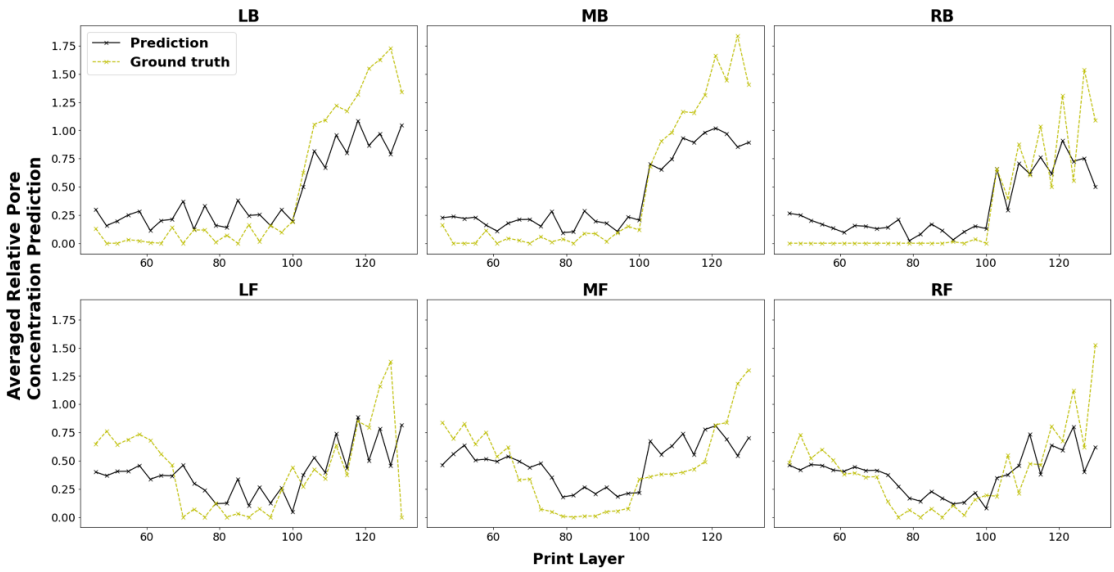


Figure 15. Test set predictions for relative pore concentration, averaged per layer and compared against ground truth.

The correlation coefficient between the averaged pore uniformity prediction and the ground truth values are overall high, with the same exception of the LF cylinder. The average NMAE value here is 0.16, thus also showing moderate to strong predictive capabilities.

4.2.2. Ablation study

An ablation study was performed to verify the impact of incorporating auxiliary features (via the AWN) and temporal information (via a temporal sliding window). Two additional models were trained using identical procedures and data: a model using only single VIS-SWIR image pairs, and a model adding auxiliary features (AWN) but lacking temporal history. Table 4 shows the averaged performance metrics of these models in comparison to the full model. It is evident that the addition of each aspect improved the results.

4.2.3. Inference time

The proposed system is implemented and optimized for execution on GPU using the Quasar programming language [64] to ensure it meets the high-speed requirements of the SLS process. The entire algorithm is implemented as a persistent GPU kernel, where tasks are spatially distributed across the

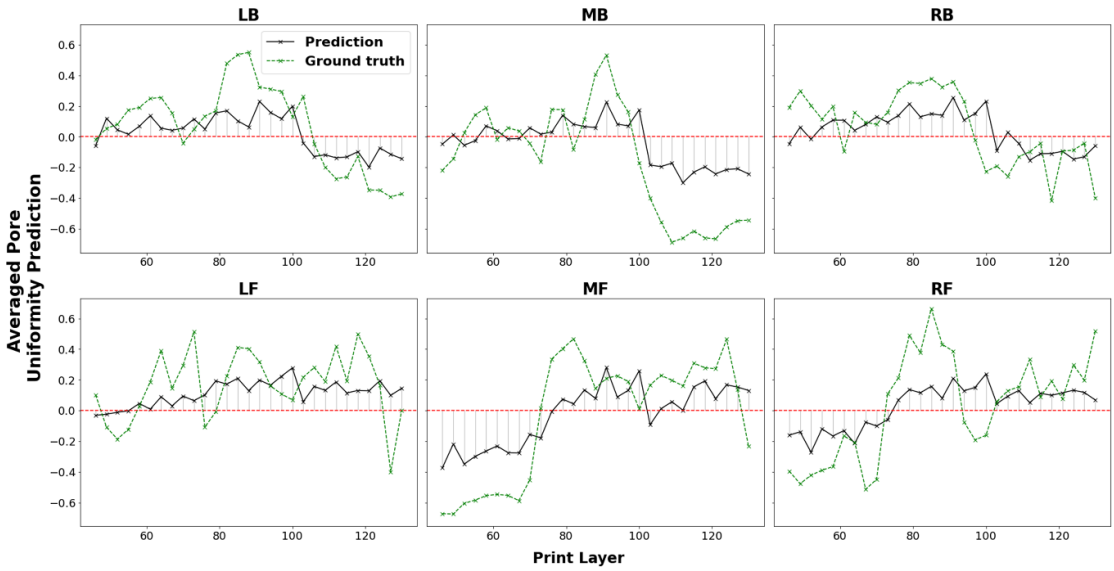


Figure 16. Test set predictions for pore uniformity, averaged per layer and compared against ground truth.

Table 2. Quantitative metrics of the predicted relative pore concentration values.

	LB	MB	RB	LF	MF	RF	AVG
$R^2 \uparrow$	0.71	0.75	0.72	0.23	0.44	0.55	0.56
NMAE $\downarrow$	0.16	0.15	0.13	0.15	0.14	0.11	0.14

Table 3. Quantitative metrics of the predicted pore uniformity values.

	LB	MB	RB	LF	MF	RF	AVG
$\rho \uparrow$	0.79	0.89	0.58	0.31	0.83	0.71	0.68
NMAE $\downarrow$	0.14	0.18	0.13	0.14	0.19	0.17	0.16

Table 4. Comparison of averaged quantitative metrics from the ablation study for (a) the relative pore concentration target and (b) the pore uniformity target. Bold values indicate the best performance for each evaluation metric.

	Images	Images + Auxiliary	Full
		(a) Concentration	
$R^2 \uparrow$	0.11	0.39	**0.56**
NMAE $\downarrow$	0.20	0.17	**0.14**
		(b) Uniformity	
$\rho \uparrow$	0.59	0.61	**0.68**
NMAE $\downarrow$	0.18	0.17	**0.16**

streaming multiprocessors. For performance evaluation, execution times were measured using a dataset of 26 000 frames captured during the printing process. The tests were conducted on the edge computing unit, using the integrated NVIDIA RTX 4060 Ti GPU. The measured execution times were 0.4 ms per frame, confirming the system's capability to operate at the high speeds required for SLS monitoring on a consumer-grade GPU.

5. Discussion

This study had two main objectives. The first objective was to define control-actionable porosity targets that can serve as ground truth targets for *in-situ* monitoring systems. These targets had to capture both the magnitude of porosity creation, and an non-ambiguous indication of its cause. The second objective

was the design of an *in-situ* monitoring system that predicts the proposed targets using high-speed cameras. The monitoring system should be able to handle fast throughput to capture the high-speed behavior of the SLS process and achieve high predictive accuracy to enable in-process control.

The underlying justification of the proposed porosity targets was validated by analyzing the patterns of porosity in multiple printed cylinders with varying laser powers. Experimental results demonstrate that increases in pore concentrations have strong correlations to off-nominal printing regimes (figure 10). Additionally, given that pore concentrations are strongly related to mechanical properties, the metric forms a robust basis for a porosity target. The experiment also confirmed the hypothesized pattern between pore concentrations and part density (figure 12), within a single print layer, in below-nominal printing zones. Physically, the relation makes sense, as insufficient melting will mainly generate porosity at the layer boundaries. It is therefore expected that the relationship will persist in any arbitrary printed part, making it another good basis for a porosity target.

Both porosity targets largely aligned with experimental expectations (figures 11–14). Some exceptions occurred for the pore uniformity measurements, primarily in transitional regions and the below-nominal zone of the LF cylinder. It is unknown why the below-nominal zone of the LF cylinder is an outlier for a large part of its layers, but suspected that the part's location in the build chamber plays a significant role. For example, in [62] it was shown that laser incidence had a large influence on energy density and porosity. In [3] and [30] it was shown that the powder bed suffers from non-homogeneous temperature distributions due to non-uniform heating elements, convection channels, and chamber wall cooling, which may lead to porosity. Additionally, the inter-layer time is different for each cylinder, which is known to influence porosity [65]. Considering that, it is hard to say for the LF cylinder if the proposed targets are not representative, or if the LF cylinder is simply deviating from the expected pattern due to other reasons.

A key strength of the porosity targets is their extraction from any arbitrary part, eliminating the need design for specific print geometries to capture accurate ground truth values. This strength allows for any printed object to be used as a training data set for data-driven *in-situ* systems. However, there is some nuance to be mentioned. First, the measurement requires high-resolution x-ray CT scanners sufficient to resolve porosity patterns within a single print layer. Based on this work, a resolution of 1/10th of the print layer thickness is sufficient. Fortunately, post-process CT analysis is well established in AM, and access to such equipment is common among professional practitioners. Second, the inherent noisiness of porosity patterns in SLS [21, 66] is causing considerable noise in the pore uniformity measure, as was shown in figure 13. Consequently, identifying expected patterns required averaging across a number of layers, which may prove challenging for smaller volumes that span fewer printing layers. However, the deviation from the nominal is rather small in these experiments ($\pm$10%). It is expected that for larger deviations, the distinction becomes clearer, and the targets appear less noisy.

The accomplishment of the second objective was evidenced by evaluating the *in-situ* monitoring system on unseen test data. Both predicted targets showed good correlations with the ground truth values and moderate to strong predictive capabilities on the single instance level, based on the NMAE values. The ablation study showed that the addition of auxiliary and temporal information strengthened the predictive capabilities of the model. Finally, the test on the inference speed of the model showed that the system can be leveraged for high-speed ($>$ 2 kHz) detection and control.

It is worth noting the variations in predictive capabilities across individual cylinders. Again, the LF cylinder appears to be an outlier with the lowest correlations for both prediction targets. This can probably be attributed to the noisiness and uncertainty of its ground truth targets. Nevertheless, the overall expected pattern in the predictions still appears. There is also a notable difference in evaluation scores between the front-row and back-row cylinders. In the front row cylinders, two transitions in printing regimes occur: below-nominal to nominal, to above-nominal. The back row cylinders only experience a nominal to below-nominal transition. Just as with the porosity targets, it is plausible that these transition regions cause different thermal behavior, making predictions more difficult based on the more limited training data.

Direct quantitative comparisons with other works are difficult due to our unique imaging setup and prediction targets. However, parallels can be drawn with the work of Hofmann *et al* [35], whose solidity predictions are analogous to our relative pore concentration. Hoffman *et al* reported an NMAE of 0.0336, yet these predictions rely on accumulating 3000 images across an entire print layer cycle (preheating, cool-down, and recoating). Their image sequence is large enough to track the entire thermal cycle of a single layer. Hence, it is expected that their predictions are more accurate. However, this limits its predictions to a layer-by-layer basis. Our contribution uses only single VIS-SWIR image pairs (and 3 previous encoded pairs) collected during laser exposure at 2000 fps, which allows for within-layer control. These image pairs only provide a momentary snapshot of the process conditions, containing more

limited information. Under these circumstances, our NMAE of 0.14 is not a large reduction in accuracy from the 0.0336 reported by Hofmann *et al* In comparison, our error is increased by a factor of about 4, but our temporal resolution is higher by a factor of 7500 (counting 4 image pairs per prediction). Additionally, our proposed system predicts multiple porosity targets to allow for control, a more complex task.

The path towards higher accuracies is clear by analyzing why works like Hofmann *et al* [35] achieve such good results: information is leveraged across the entire thermal cycle. Low-speed sensors such as long-wave infrared imagers could complement our high-speed system by monitoring powder bed temperatures across key phases (pre-recoating, post-recoating, pre-sintering). Additionally, other sensors could measure the temperature of the build chamber, the heating elements, etc. In summary, a hybrid system that incorporates low-speed sensors can complement the information captured from the high-speed sensors to improve predictions. In our proposed system, this information can easily be integrated as additional auxiliary features.

Overall, this study offers key insights into SLS porosity formation and its *in-situ* monitoring. The validation of the dual porosity targets (concentration and uniformity) confirms that porosity arising from different energy input regimes exhibits distinct spatial signatures within a layer, providing a quantifiable link between process deviations and defect morphology. Furthermore, this work demonstrates that synchronized visual and SWIR imaging captures critical dynamics for porosity development in SLS. The monitoring system's performance, particularly the improvements shown by incorporating contextual process information (via the AWN) and temporal data (via buffering), highlights the importance of considering both local melt pool appearance and broader process conditions for accurate prediction. Crucially, these predictions are directly suited for guiding process control. The dual numeric outputs provide actionable insights, and their potentially smooth evolution offers advantages for stable, gradual control adjustments: e.g. if predicted relative pore concentration rises while uniformity is negative (indicating likely insufficient melting), a control system could locally increase the delivered energy density (e.g. by adjusting laser power). Conversely, high concentration with positive uniformity might suggest overheating, potentially triggering a local energy density reduction.

The system's high-speed (2 kHz) generation of control-relevant signals paves the way for implementing advanced intra-layer feedback mechanisms. The system operates with a cycle time of 0.5 ms, which accommodates the measured inference latency of 0.4 ms. Even at the high scanning speeds used in this study (4.4 ms^{-1}), this corresponds to a spatial resolution of approximately 2.2 mm. This technical capability suggests the feasibility of effective vector-by-vector or even intra-vector power modulation via high-frequency controllers, such as the MCP. While this level of control necessitates high-speed interfaces, systems without such interfaces can still benefit by aggregating the high-speed predictions to trigger layer-by-layer adjustments, such as dynamically updating the build file for the subsequent layer. Industrially, the further development of these capabilities could facilitate a shift from costly post-process inspections to active in-process quality assurance [67]. By mitigating defects as they occur, such systems offer the potential to reduce scrap rates and validation costs, addressing two of the primary economic barriers to the broader adoption of SLS for critical end-use parts.

Finally, we acknowledge several limitations shaping the scope of this study. The experiments utilized simple cylindrical geometries in order to isolate the effects of energy density to the laser power settings, potentially limiting immediate generalization to complex shapes. Crucially, however, the system's architecture explicitly accounts for local geometric variations through the inclusion of geometric auxiliary inputs via the AWN. Our previous works on applying AWNs in SLM [39, 57] have successfully demonstrated generalization to complex geometric features such as overhangs and thin walls. Consequently, while the proof-of-concept shown in this study was simple, the architectural foundation for handling complex, industrially relevant shapes is both present and capable of generalization. Regarding process parameters, variations were restricted to only three distinct laser power configurations (nominal, +10% and –10%). It is well established that other parameters, such as scan speed and hatch distance, also significantly influence porosity formation. Thus, future work should include a wider range of process parameters, and more parameter configurations to provide extra evidence and better quantify the sensitivity of the system. Validation across diverse SLS polymers (e.g. PA11, Polypropylene) is also needed. Applying this methodology to new materials will necessitate specific adaptations, particularly algorithmic retraining, due to differing optical and thermal properties. It must also be recognized that the CT resolution limits the ground truth to pores roughly 10 μm or larger, potentially excluding finer porosity known to occur in SLS [68]. While the physical basis of the porosity targets suggests broader relevance across conditions, and since other parameters ultimately modulate the same energy balance, it is plausible that deviations from their optimal values could induce similar types of porosity phenomena to those captured in this study. Still, extensive experimental verification of the monitoring system across these expanded domains is still required. Nevertheless, with this work, we have shown that control-actionable

porosity targets can be defined, and accurate high-speed *in-situ* monitoring systems can be designed, to predict, and more importantly, enable in-process control for porosity in the SLS process.

6. Conclusion

This study advances porosity control in SLS by defining physically grounded, actionable porosity targets and demonstrating their predictability using a high-speed *in-situ* monitoring system. By relying on pore concentration and spatial pore uniformity measures derived from standard x-ray CT data, the proposed targets quantify both the degree of pore formation and the appropriate corrective actions to prevent it. Validated through experiments with deliberately altered print parameters, these targets consistently represent off-nominal conditions and do not require custom experimental designs, allowing any printed part to contribute to training and refining *in-situ* monitoring capabilities.

Leveraging high-speed visible and SWIR imaging, a monitoring system was developed using context-adaptive neural networks and a temporal sliding window to accurately predict the proposed porosity targets at rates exceeding 2 kHz. The system achieved normalized mean absolute error values of 14% for predicting relative pore concentrations and a correlation coefficient of 0.68 for predicting the cause of the porosity on an unseen test set. Ultimately, this approach may reduce reliance on costly post-process inspections, enhance part quality, and open new avenues for closed-loop control in polymer SLS.

Data availability statement

The data that support the findings of this study are openly available at the following URL/DOI: http://doi.org/10.17 605/OSF.IO/GHCVJ.

Acknowledgments

This work was financially supported by the imec ICON VLAIO project 'MultipLICITY' (HBC.2022.0094), a collaboration between TELIN-IPI-imec, Ghent University; Flanders MAKE@KULeuven; Materialise; imec Vision Lab, University of Antwerp; Flanders MAKE; Dekimo; Manufacturing Processes and Systems (MaPS); Xenics; Rejig; and AdditiveLab. This research received funding from the Flemish Government under the 'Onderzoeksprogramma Artificiele Intelligentie (AI) Vlaanderen' programme.

Conflict of interest

No potential conflict of interest was reported by the author(s).

Author contributions

Mathieu Vandecasteele 0000-0001-8346-131X
Conceptualization (lead), Formal analysis (lead), Investigation (lead), Methodology (lead), Software (lead), Validation (lead), Visualization (lead), Writing – original draft (lead), Writing – review & editing (lead)

Ricardo Santander
Investigation (supporting), Methodology (supporting)

Domenico Iuso 0000-0001-7919-7878
Investigation (supporting)

Mohsen Nourazar 0000-0001-9733-0476
Investigation (supporting), Methodology (supporting), Software (supporting)

Wilfried Philips
Writing – review & editing (supporting)

Sven Cornelissen
Conceptualization (supporting), Funding acquisition (equal), Project administration (equal)

Brian G Booth 0000-0001-9228-4810
Conceptualization (supporting), Funding acquisition (equal), Methodology (supporting), Project administration (equal), Software (supporting), Supervision (lead), Validation (supporting), Writing – review & editing (supporting)

References

[1] Kafle A, Luis E, Silwal R, Pan H M, Shrestha P L and Bastola A K 2021 3D/4D printing of polymers: fused deposition modelling (FDM), selective laser sintering (SLS) and stereolithography (SLA) *Polymers* **13** 3101
[2] Morano C and Pagnotta L 2023 Additive manufactured parts produced by selective laser sintering technology: porosity formation mechanisms *J. Polym. Eng.* **43** 537–55
[3] Bourell D L, Watt T J, Leigh D K and Fulcher B 2014 Performance limitations in polymer laser sintering *Phys. Proc.* **56** 147–56
[4] Dupin S, Lame O, Barrés C and Charmeau J-Y 2012 Microstructural origin of physical and mechanical properties of polyamide 12 processed by laser sintering *Eur. Polym. J.* **48** 1611–21
[5] Rosso S, Meneghello R, Biasetto L, Grigolato L, Concheri G and Savio G 2020 In-depth comparison of polyamide 12 parts manufactured by Multi Jet fusion and selective laser sintering *Addit. Manuf.* **36** 101713
[6] Schmid M 2018 *Laser Sintering With Plastics* (Hanser) (https://doi.org/10.3139/9781569906842.fm)
[7] Yuan S, Li J, Yao X, Zhu J, Gu X, Gao T, Xu Y and Zhang W 2020 Intelligent optimization system for powder bed fusion of processable thermoplastics *Addit. Manuf.* **34** 101182
[8] Peyre P, Rouchausse Y, Defauchy D and Régnier G 2015 Experimental and numerical analysis of the selective laser sintering (SLS) of PA12 and PEKK semi-crystalline polymers *J. Mater. Process. Technol.* **225** 326–36
[9] Ryse R V, Edeleva M, Van Stichel O, D'hooge D R, Pille F, Fiorio R, De Baets P and Cardon L 2022 Setting the optimal laser power for sustainable powder bed fusion processing of elastomeric polyesters: a combined experimental and theoretical study *Materials* **15** 385
[10] Nar K, Majewski C and Lewis R 2022 A comprehensive characterisation of Laser Sintered Polyamide-12 surfaces *Polym. Test.* **106** 107450
[11] Ho H C H, Gibson I and Cheung W L 1999 Effects of energy density on morphology and properties of selective laser sintered polycarbonate *J. Mater. Process. Technol.* **89-90** 204–10
[12] Stichel T, Frick T, Laumer T, Tenner F, Hausotte T, Merklein M and Schmidt M 2017 A Round Robin study for Selective Laser Sintering of polyamide 12: microstructural origin of the mechanical properties *Opt. Laser Technol.* **89** 31–40
[13] Calignano F, Giuffrida F and Galati M 2021 Effect of the build orientation on the mechanical performance of polymeric parts produced by multi jet fusion and selective laser sintering *J. Manuf. Process.* **65** 271–82
[14] Liebrich A, Langowski H C, Schreiber R and Pinzer B R 2019 Porosity distribution in laser-sintered polymeric thin sheets as revealed by X-ray micro tomography *Polym. Test.* **76** 286–97
[15] Pavan M, Craeghs T, Verhelst R, Ducatteeuw O, Kruth J-P and Dewulf W 2016 CT-based quality control of Laser Sintering of Polymers *Case Stud. Nondestruct. Test. Eval.* **6** 62–68
[16] Ling Z, Wu J, Wang X, Li X and Zheng J 2019 Experimental study on the variance of mechanical properties of polyamide 6 during multi-layer sintering process in selective laser sintering *Int. J. Adv. Manuf. Technol.* **101** 1227–34
[17] Yang F, Jiang T, Lalier G, Bartolone J and Chen X 2020 A process control and interlayer heating approach to reuse polyamide 12 powders and create parts with improved mechanical properties in selective laser sintering *J. Manuf. Process.* **57** 828–46
[18] Yang F and Chen X 2021 A combined theoretical and experimental approach to model polyamide 12 degradation in selective laser sintering additive manufacturing *J. Manuf. Process.* **70** 271–89
[19] Allison J, Sharpe C and Seepersad C C 2019 Powder bed fusion metrology for additive manufacturing design guidance *Additive Manuf.* **25** 239–51
[20] Lu Q Y and Wong C H 2017 Applications of non-destructive testing techniques for post-process control of additively manufactured parts *Virtual Phys. Prototyping* **12** 301–21
[21] Dewulf W, Pavan M, Craeghs T and Kruth J-P 2016 Using x-ray computed tomography to improve the porosity level of polyamide-12 laser sintered parts *CIRP Ann.* **65** 205–8
[22] Kanchi S P, Sundari S, Rao C V S P and Raju S 2024 Application of response surface methodology and machine learning approaches to predict porosity and compression strength of selective laser-sintered polyamide-12 scaffold parts *J. Mater. Eng. Perform.* **34** 6274–83
[23] Li M, Han Y, Zhou M, Chen P, Gao H, Zhang Y and Zhou H 2020 Experimental investigating and numerical simulations of the thermal behavior and process optimization for selective laser sintering of PA6 *J. Manuf. Process.* **56** 271–9
[24] Yeh H-P, Pan Z, Bayat M and Hattel J H 2024 A thermal fluid dynamic model for the melt region during the laser powder bed fusion of polyamide 11 (PA11) *J. Manuf. Process.* **124** 1422–37
[25] Moges T, Ameta G and Witherell P 2019 A review of model inaccuracy and parameter uncertainty in laser powder bed fusion models and simulations *J. Manuf. Sci. Eng.* **141** 040801
[26] Tao Z, Sinico M, Vrancken B and Dewulf W 2025 High-fidelity surrogate modelling for geometric deviation prediction in laser powder bed fusion using in-process monitoring data *Virtual Phys. Prototyping* **20** e2523550
[27] Safdar M, Wood G, Zimmermann M, Lamouche G, Wanjara P and Zhao Y F 2025 Detecting the extent of co-existing anomalies in additively manufactured metal matrix composites through explainable selection and fusion of multi-camera deep learning features *Virtual Phys. Prototyping* **20** e2515240
[28] Akmal J, Minet-Lallemand K, Kuva J, Syvänen T, Ylander P, Puttonen T, Björkstrand R, Partanen J, Nyrhilä O and Salmi M 2025 AI-based defect detection and self-healing in metal additive manufacturing *Virtual Phys. Prototyping* **20** e2500671
[29] Zhao J, Yang Z, Chen Q, Zhang C, Zhao J, Zhang G, Dong F and Liu S 2025 Real-time detection of powder bed defects in laser powder bed fusion using deep learning on 3D point clouds *Virtual Phys. Prototyping* **20** e2449171
[30] Phillips T, Fish S and Beaman J 2018 Development of an automated laser control system for improving temperature uniformity and controlling component strength in selective laser sintering *Addit. Manuf.* **24** 316–22
[31] Phillips T, Ricker T, Fish S and Beaman J 2020 Design of a laser control system with continuously variable power and its application in additive manufacturing *Addit. Manuf.* **34** 101173
[32] Williams R J, Piglione A, Rønneberg T, Jones C, Pham M-S, Davies C M and Hooper P A 2019 In situ thermography for laser powder bed fusion: Effects of layer temperature on porosity, microstructure and mechanical properties *Addit. Manuf.* **30** 100880
[33] Guan G, Hirsch M, Syam W P, Leach R K, Huang Z and Clare A T 2016 Loose powder detection and surface characterization in selective laser sintering via optical coherence tomography *Proc. R. Soc. A* **472** 20160201
[34] Schlicht S, Jaksch A and Drummer D 2022 Inline quality control through optical deep learning-based porosity determination for powder bed fusion of polymers *Polymers* **14** 885
[35] Hofmann J, Li Z, Taphorn K, Herzen J and Wudy K 2024 Porosity prediction in laser-based powder bed fusion of polyamide 12 using infrared thermography and machine learning *Addit. Manuf.* **85** 104176

[36] Meier C, Weissbach R, Weinberg J, Wall W A and Hart A J 2019 Critical influences of particle size and adhesion on the powder layer uniformity in metal additive manufacturing *J. Mater. Process. Technol.* **266** 484–501
[37] Greiner S, Wudy K, Wörz A and Drummer D 2019 Thermographic investigation of laser-induced temperature fields in selective laser beam melting of polymers *Opt. Laser Technol.* **109** 569–76
[38] Ahar A, Vandecasteele M, Booth B G, Grave K D, Verhees D, Philips W and Bey-Temsamani A 2024 Combining high-throughput imaging in visible and SWIR wavelengths for *in-situ* porosity prediction in laser powder bed fusion *Proc. SPIE* **12876** 46–58
[39] Vandecasteele M, Iuso D, Nasab M H, Verhees D, Sanctorum J, Nourazar M, Bey-Temsamani A and Booth B G 2024 High-speed deformation prediction in selective laser melting using context-adaptive neural networks *2024 Int. Conf. on Machine Learning and Applications (ICMLA)* pp 1712–9
[40] Hooper P A 2018 Melt pool temperature and cooling rates in laser powder bed fusion *Addit. Manuf.* **22** 548–59
[41] Haeri S, Wang Y, Ghita O and Sun J 2017 Discrete element simulation and experimental study of powder spreading process in additive manufacturing *Powder Technol.* **306** 45–54
[42] Beitz S, Uerlich R, Bokelmann T, Diener A, Vietor T and Kwade A 2019 Influence of powder deposition on powder bed and specimen properties *Materials* **12** 297
[43] Schmid M and Wegener K 2016 Additive manufacturing: polymers applicable for laser sintering (LS) *Proc. Eng.* **149** 457–64
[44] Vasquez M, Haworth B and Hopkinson N 2011 Optimum sintering region for laser sintered nylon-12 *Proc. Inst. Mech. Eng. B* **225** 2240–8
[45] Samber B D, Renders J, Elberfeld T, Maris Y, Sanctorum J, Six N, Liang Z, Beenhouwer J D and Sijbers J 2021 FleXCT: a flexible X-ray CT scanner with 10 degrees of freedom *Opt. Express* **29** 3438–57
[46] Vandecasteele M, Heylen R, Iuso D, Thanki A, Philips W, Witvrouw A, Verhees D and Booth B G 2023 Towards material and process agnostic features for the classification of pore types in metal additive manufacturing *Mater. Des.* **227** 111757
[47] Stichel T, Frick T, Laumer T, Tenner F, Hausotte T, Merklein M and Schmidt M 2018 A Round Robin study for selective laser sintering of polymers: back tracing of the pore morphology to the process parameters *J. Mater. Process. Technol.* **252** 537–45
[48] Moeskops E, Kamperman N, van de Vorst B and Knoppers R 2004 Creep behaviour of polyamide in selective laser sintering *2004 Int. Solid Freeform Fabrication Symp. (SFF)* (https://doi.org/10.26153/tsw/5708)
[49] Bugeda G, Cervera M and Lombera G 1999 Numerical prediction of temperature and density distributions in selective laser sintering processes *Rapid Prototyping J.* **5** 12–26
[50] Yamauchi Y, Kigure T and Niino T 2023 Penetration depth optimization for proper interlayer adhesion using near-infrared laser in a low-temperature process of PBF-LB/P *J. Manuf. Process.* **98** 126–37
[51] Pavan M, Craeghs T, Kruth J-P and Dewulf W 2016 Understanding the laser sintering of polymers at microscale level by using x-ray computed tomography *2016 Int. Solid Freeform Fabrication Symp. (SFF)* (available at: https://hdl.handle.net/2152/89660)
[52] Sillani F, MacDonald E, Villela J, Schmid M and Wegener K 2022 *in-situ* monitoring of powder bed fusion of polymers using laser profilometry *Addit. Manuf.* **59** 103074
[53] Chen Y and Medioni G 1992 Object modelling by registration of multiple range images *Image Vis. Comput.* **10** 145–55
[54] Vandecasteele M, Verhees D, Philips W, Bey-Temsamani A and Booth B G 2024 High-speed interpretable representation learning of melt pool signatures in metal additive manufacturing *2024 IEEE 20th Int. Conf. on Automation Science and Engineering (CASE)* pp 1618–24
[55] Higgins I, Matthey L, Pal A, Burgess C P, Glorot X, Botvinick M, Mohamed S and Lerchner A 2017 beta-VAE: learning basic visual concepts with a constrained variational framework *Int. Conf. on Learning Representations (ICLR)* (available at: https://openreview.net/forum?id = Sy2fzU9gl)
[56] Shao H, Yao S, Sun D, Zhang A, Liu S, Liu D, Wang J and Abdelzaher T 2020 Controlvae: controllable variational autoencoder *Proc. 37th Int. Conf. on Machine Learning (ICML)* (https://doi.org/10.48550/arXiv.2004.05988)
[57] Vandecasteele M, searle S, Iuso D, Nourazar M, Ahar A, Nasab M H, Vrancken B and Booth B 2025 A high-speed multimodal *in-situ* monitoring system for detecting porosity and deformation defects in the powder bed fusion of metallic parts *Euro PM2025 Proc., European Powder Metallury Association (EPMA)* p 9
[58] Pei E, Jiang D and Sahli H 2020 An efficient model-level fusion approach for continuous affect recognition from audiovisual signals *Neurocomputing* **376** 42–53
[59] Soldner D, Greiner S, Burkhardt C, Drummer D, Steinmann P and Mergheim J 2021 Numerical and experimental investigation of the isothermal assumption in selective laser sintering of PA12 *Addit. Manuf.* **37** 101676
[60] Lapointe S, Guss G, Reese Z, Strantza M, Matthews M J and Druzgalski C L 2022 Photodiode-based machine learning for optimization of laser powder bed fusion parameters in complex geometries *Addit. Manuf.* **53** 102687
[61] Kingma D P and Ba J 2015 Adam: a method for stochastic optimization *3rd Int. Conf. on Learning Representations (ICLR)* (https://doi.org/10.48550/arXiv.1412.6980)
[62] Fathi-Hafshejani P, Soltani-Tehrani A, Shamsaei N and Mahjouri-Samani M 2022 Laser incidence angle influence on energy density variations, surface roughness and porosity of additively manufactured parts *Addit. Manuf.* **50** 102572
[63] Ng W L, Goh G L, Goh G D, Ten J S J and Yeong W Y 2024 Progress and opportunities for machine learning in materials and processes of additive manufacturing *Adv. Mater.* **36** 2310006
[64] Goossens B, De Vylder J and Philips W 2014 Quasar - a new heterogeneous programming framework for image and video processing algorithms on CPU and GPU *2014 IEEE Int. Conf. on Image Processing (ICIP)* (https://doi.org/10.1109/ICIP.2014.7025441) pp 2183–5
[65] Pavan M, Faes M, Strobbe D, Van Hooreweder B, Craeghs T, Moens D and Dewulf W 2017 On the influence of inter-layer time and energy density on selected critical-to-quality properties of PA12 parts produced via laser sintering *Polym. Test.* **61** 386–95
[66] Pavan M 2018 CT-based quality optimization of laser sintering of polyamide-12 *PhD Thesis* (KULeuven) (available at: https://lirias.kuleuven.be/1674951&lang = en)
[67] Chen L, Bi G, Yao X, Su J, Tan C, Feng W, Benakis M, Chew Y and Moon S K 2024 *in-situ* process monitoring and adaptive quality enhancement in laser additive manufacturing: a critical review *J. Manuf. Syst.* **74** 527–74
[68] Rouholamin D and Hopkinson N 2014 An investigation on the suitability of micro-computed tomography as a non-destructive technique to assess the morphology of laser sintered nylon 12 parts *Proc. Inst. Mech. Eng. B* **228** 1529–42