

Full length article

Adaptive In-Situ monitoring for laser powder bed Fusion: Self-Supervised learning for layer thickness monitoring across scan lengths based on pyrometry

Bariş Kavas^a, Roland Axel Richter^b, Michael R. Tucker^a, Vigneashwara Pandiyan^{c,*}

^a Advanced Manufacturing Laboratory, Department of Mechanical and Process Engineering, ETH Zürich, Zürich 8092, Switzerland

^b Laboratory for Advanced Materials Processing (LAMP), Swiss Federal Laboratories for Materials Science and Technology (Empa) CH-3602 Thun, Switzerland

^c Digital Manufacturing and Surface Engineering (DMS), Department of Mechanical and Materials Engineering, University of Turku FI-20014, Turun yliopisto, Finland

ARTICLE INFO

Keywords:

Laser Powder Bed Fusion
Process Monitoring
In-Situ Sensing
Optical Emission
Temporal Convolutional Network

ABSTRACT

Laser Powder Bed Fusion (LPBF) is a widely used additive manufacturing process that offers high precision and design flexibility but suffers from quality inconsistencies due to variations in layer thickness. Ensuring uniform layer thickness is critical, as deviations can lead to defects such as porosity and geometric distortion. Existing inspection methods rely on optical or thermographic imaging techniques that limit spatial and temporal resolution and require supervised machine learning techniques. This study introduces a novel self-supervised machine learning approach leveraging on-axis pyrometry data to infer local layer thickness variations during LPBF. A Temporal Convolutional Network (TCN) is trained using a unique data randomization technique to handle variable-length time-series signals. The model is designed to learn representations without requiring labelled data, addressing a key challenge in real-time process monitoring. Experimental validation was conducted using a controlled LPBF setup with varying layer thicknesses. The trained model successfully classified different thickness regimes and demonstrated the ability to capture process anomalies such as short-feeding or warping. Analysis using t-distributed stochastic neighbour embedding (t-SNE) revealed well-separated clusters for distinct layer thicknesses, validating the model's effectiveness. However, the sensor's resolution limited discrimination below 20 μm , highlighting the need for sensor fusion strategies. Future work will focus on integrating additional data sources, such as acoustic emissions and photonic sensing, to improve resolution and extend the model's applicability to complex geometries and scan patterns. The proposed method provides a foundation for real-time LPBF quality control, enabling adaptive process optimization and defect prevention, paving the way for industrial-scale adoption of in-situ monitoring solutions.

1. Introduction

Laser powder bed fusion (LPBF) additive manufacturing (AM) is an ideal production technique for prototype development end-use parts with complex designs and high-mix, low-volume production runs. Compared to other metal AM modalities with similar build volumes, it offers greater design freedom and spatial resolution [1]. While special design procedures and geometries such as lattice structures, internal channels, and topographic optimization are made feasible by LPBF, parts with conventional designs can also be manufactured more rapidly and with less waste for prototyping purposes [2]. Despite its apparent simplicity, the LPBF process involves complex melt pool dynamics characterized by rapid heating and cooling interactions between the

laser beam and powder, the characteristics of the powder material, environmental conditions, and repeated thermal cycles, making industrial adoption challenging [3]. The variable quality of material processing in LPBF has slowed its adoption in producing components for safety-critical applications [4,5]. Although there is robust demand for LPBF, its inconsistent repeatability poses a significant barrier to broader industrial use [6]. The quality of the components during operation heavily relies on the uniformity and density of the scanned layers. Subsequent layers in the LPBF process can mask defects within a layer, adversely impacting the part's fatigue life and strength due to its layer-wise construction. To meet industrial standards and specifications, metal-based additively manufactured components must be shown to have acceptable levels of critical defects during and after production [7].

* Corresponding author.

E-mail address: vigneashwara.solairajapandiyan@utu.fi (V. Pandiyan).

<https://doi.org/10.1016/j.optlastec.2025.114070>

Received 27 March 2025; Received in revised form 22 August 2025; Accepted 4 October 2025

Available online 16 October 2025

0030-3992/© 2025 The Author(s). Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

In LPBF, a laser melts a thin layer of powdered material, rapidly solidifying to form a weld bead [8]. Parts are fabricated by stacking overlapping weld beads in the horizontal and vertical directions until a part is formed. Therefore, information from the entire part volume can be acquired by observing the melt pool characteristics throughout the process. This observability provides an advantage for LPBF over conventional manufacturing processes since it may allow direct in-situ qualification and process control. Ex-situ inspection methods are commonly used for LPBF-manufactured parts today, which can only identify defects after production, resulting in wasted time, energy, and money for machine operation, materials, and labour on defective parts [9,10]. In-situ process monitoring leverages the interaction between the laser and powder, generating heat, melting, and vaporizing the material, producing detectable optical, acoustic, and electromagnetic emissions [11,12]. Literature on various techniques for LPBF is abundant over a wide range of temporal and spatial scales, with comprehensive reviews available on the subject [13,14]. Optical emission-based techniques are widely used for non-contact thermal measurements by capturing radiation from a heated surface. Among these, Pyrometry is a well-established method that quantifies infrared radiation intensity to determine temperature, making it a valuable tool for high-speed process monitoring in LPBF. Among on-axis process monitoring techniques, Pyrometry is distinguished by its ability to identify various process states and defects at a high sampling rate [15,16]. Pyrometers are commonly used in monitoring systems to measure thermal emissions by quantifying the intensity of infrared radiation emitted by an object, thereby allowing non-contact measurement. In the LPBF process, the black body radiation emitted by the melt pool as a high-temperature medium is captured, typically with sampling rates of 25–100 kHz [17]. Pyrometers have been studied extensively in the literature and applied to identify process anomalies such as keyholing, lack-of-fusion, and local overheating [18–22]. Besides detecting defects, on-axis pyrometers are employed for part-scale digital twin generation for post-process investigation of the part quality [23]. However, despite their widespread application, no studies were found investigating the feasibility of using pyrometry to evaluate local powder layer thickness variation, a common cause of LPBF defects.

The layer height (i.e., the distance descended by the build plate for each layer) is one of the most critical parameters identified and fixed early during process development [24]. More precisely, it is critical to stabilize the effective layer thickness, which is the local thickness of the powder layer over the underlying part or substrate to be consolidated [25]. For brevity, the term layer thickness is used here throughout. Although the process parameters are optimized for a specific layer thickness, local or layer-wise thickness may change during the process due to various anomalies such as powder short-feeding [26], layer delamination [27], excessive distortion [28], support structure failure [29], or a damaged recoater blade [30,31]. In summary, deviation in the thickness of the layer is a serious threat to part quality. A region with decreased layer thickness is susceptible to overheating and thermal distortion due to increased energy absorption by the consolidated material [32,33]. In contrast, an increased layer thickness is prone to lack-of-fusion porosity [34,35]. Therefore, the in-situ identification of the layer thickness deviation is very important to ensure part quality. Several studies in the literature focus on the in-situ identification of layer thickness. Williams et al. proposed a method with a laser displacement sensor to measure the actual layer thickness during the build [36]. Their work achieved accurate measurements; however, it requires specialized measurement hardware, and data acquisition delays the overall build time. Liu et al. leveraged an off-axis thermal camera to acquire a surface temperature signal after the recoating and correlated the signal change with the layer thicknesses [37,38]. Although the proposed method is shown to provide a good estimation, the spatial resolution of their method is limited by the resolution of the off-axis thermal camera. There are no studies in the literature that leverages the on-axis pyrometer data to identify the layer thickness. Inference of

the layer thickness value by analyzing the optical emission data potentially enables the estimation and prevention of the aforementioned process defects. While the part quality can be evaluated through the estimation, prevention can be performed by closed-loop control during the process [39]. However, the evaluation of the pyrometer signal poses significant challenges due to its one-dimensional nature. Variations in layer thickness have been directly linked to the formation of defects such as cracks and porosity, which compromise part integrity and reliability [40,41]. Addressing this challenge through in-situ monitoring is therefore critical for advancing qualification of LPBF builds.

1.1. Model-based signal processing

Understanding the complex, non-linear temporal and spatial data generated by sensors during the LPBF process is difficult for human operators due to the nature of laser-material interactions [42–45]. Using signal processing techniques, statistical features that are sparse and lower-dimensional can be extracted from the waveform data in the time, frequency, and time–frequency domains [46,47]. These statistical descriptors have proven useful in differentiating and interpreting key states of the LPBF process [48]. Building on this, researchers have utilized conventional machine learning (ML) methods, such as Decision Trees (DT) and K-nearest neighbors (KNN), to correlate these features with measures of part quality derived from optical and thermal observations [49,50]. Other work has focused on mapping pixel-level or multidimensional image features from build layers to final part integrity via linear support vector machines (SVM) [51]. Feature extraction approaches like multilinear principal component analysis (MPCA) or randomized matrix factorization (e.g., SVD) have also been explored as means of reducing data dimensionality prior to classification [52,53]. Additionally, ensemble models (e.g., Random Forest) and logistic regression have successfully leveraged signal-based features—in some cases, derived from acoustic emissions—to classify different LPBF processing regimes with reasonable reliability [47]. However, when ML solutions rely on manually engineered features, their performance is strongly tied to the features chosen, which in turn depends on the expertise of those selecting them—an inherent drawback of these methods. In response, recent research has increasingly focused on leveraging deep learning (DL) approaches for LPBF monitoring. Unlike traditional ML approaches that require data to be pre-processed into sparse representations based on human understanding, DL methods can handle raw sensor data with minimal pre-processing [54]. Deep learning, primarily via Convolutional Neural Networks (CNNs), has proven effective at detecting LPBF process anomalies like porosity, varying density levels, and delamination, analyzing 2D spectrograms, melt pool images, and IR wavelength images [55,56]. CNNs have also been used to analyze heterogeneous data streams from the process zone to predict real-time part quality [12,57,58]. Anomalies such as warpage, part shifting, and insufficient powder supply have been identified by supervised training digital images using deep residual and region proposal neural networks. Furthermore, semi-supervised training of CNNs has been advocated as a monitoring strategy that minimizes the effort required for data collection [59,60].

Recent advances in data-driven process monitoring and sensor technologies [47,61,62], have laid the groundwork for the development of Self-Supervised Learning (SSL) methods in LPBF, which this work seeks to advance. Unlike fully supervised approaches that rely heavily on large labelled datasets [63–66]. SSL leverages intrinsic structures or auxiliary tasks within unlabeled data to learn robust feature representations, eliminating costly manual labelling [67]. This methodology addresses key challenges in industrial-scale LPBF, where vast amounts of build data and part-to-part variability make it unclear which process parameters most strongly influence outcomes [68–72]. Previous studies using self-supervised CNN architectures have successfully detected melt pool anomalies like lack-of-fusion and keyhole porosity by extracting meaningful context from unannotated signals or images, and further

work has explored temporal or frequency-based SSL to handle acoustic and thermal emissions [67,73]. These self-learned embeddings promise transfer learning, enabling rapid adaptation to new process maps and improved generalizability with minimal retraining [73]. In addition, SSL circumvents the risk of contradictions introduced by human labelling, making it more flexible and straightforward to implement than traditional supervised techniques. As this approach continues to evolve, it is poised to streamline data acquisition, reduce post-mortem inspection needs, and enhance cost-efficiency in adaptive LPBF process control.

1.2. Contribution

Despite extensive research on in-situ LPBF process monitoring, significant challenges persist in effectively leveraging on-axis pyrometry for local layer thickness identification. To our knowledge, no previous work has demonstrated the feasibility of using a self-supervised Temporal Convolutional Network (TCN) to extract layer thickness information directly from pyrometer signals. This represents a key novelty of our work, as the proposed approach enables geometry-agnostic monitoring across variable scan lengths without requiring labeled datasets or post-process characterization. Based on the literature review, several research gaps are identified. First, correlations between on-axis pyrometer signals and layer thickness have primarily been studied under supervised settings and post-process part characterization. To the authors' knowledge, no prior studies have demonstrated successful real-time identification of layer thickness from in-situ pyrometer data despite the known dependence. Second, although numerous time-series classification and anomaly detection algorithms have been applied to LPBF processes with similar sensors, no documented methods account for the variation in vector lengths that arise from complex part shapes and scan strategies. Existing approaches, therefore, remain constrained by extensive labelling requirements, fixed time-series window sizes, and limited adaptability. Additionally, the inability to capture subtle layer thickness variations hinders more proactive defect prevention measures.

Considering the identified research gaps, the main contributions of this work are:

- Proposition and experimental validation of a self-supervised TCN framework for in-situ monitoring in LPBF. The model leverages high-frequency, unlabeled pyrometer signals to infer local layer thickness without requiring manual annotations or post-process characterization. This approach significantly reduces the dependency on ground truth data while preserving predictive fidelity.
- Development and integration of a novel temporal sampling algorithm designed to handle variable-length scan vectors, which are intrinsic to LPBF due to geometric complexity and toolpath diversity. This strategy facilitates robust and scalable TCN training, ensuring the model generalizes effectively across a wide range of process conditions and build geometries.
- Implementation and assessment of a real-time TCN-based monitoring framework capable of detecting deviations in layer thickness with a resolution limit of approximately 20 μm . The framework reveals the limitations of single-sensor configurations and underscores the importance of sensor fusion—integrating pyrometry with complementary modalities such as visible-spectrum or IR imaging—to achieve enhanced spatial resolution and improved defect localization in industrial LPBF systems.
- A demonstration that the TCN encoder implicitly captures melt pool physics, learning a smooth, continuous feature embedding that organizes naturally according to layer thickness, without relying on explicit physical models or governing equations.

This paper is organized into six sections. Section 1 provides an overview of the research landscape and highlights the motivation and scope of the study. Section 2 explains the experimental setup, including materials, equipment, and data acquisition methods. Section 3

introduces the proposed self-supervised learning approach and describes the core methodological contributions. Section 4 then delves into experimental results and analysis, illustrating both the efficacy and limitations of the developed framework. Section 5 summarizes the main findings and examines broader implications, while Section 6 concludes the paper by synthesizing the key insights and proposing directions for future research.

2. Experimental setup

2.1. Machine, Instrumentation, and material

The experiments were performed on Aconity3D Midi+ (Aconity3D GmbH) LPBF machine [35] equipped with a continuous wave Gaussian mode fiber laser with a wavelength of 1080 nm and a maximum output of 500 W (nLIGHT Alta), focused to a diameter ($d4\sigma$) of 80 μm . The build plate was $\varnothing 250 \text{ mm} \times 25 \text{ mm}$ thick and made of stainless steel 1.4301. A silicon recoater blade was used, and the powder used in this study was gas-atomized stainless steel 316L (1.4404) with a particle size distribution of 15 μm to 45 μm (CT POWDERRANGE 316LF, Carpenter Additive). The schematic of the on-axis implementation of the pyrometer sensor is shown in Fig. 1. The laser beam was generated and delivered to the machine's optical components via a fibre optic cable, where it was then collimated to achieve infinite focus. Next, it passed through a dynamic focusing unit that continuously adjusted to ensure consistent focus across the build platform. The beam was then deflected off a dichroic mirror into the scanner, which steered the beam over the processing area. The mirror's coating reflects the laser's specific frequency while transmitting the rest of the spectrum, allowing black body radiation from the laser exposure region to pass through. The transmitted emission is then passed through a manual focusing unit and an alignment table before reaching the pyrometer that captures the optical emission at 2 – 2.2 μm . The pyrometer captured the emitted radiation, measuring its intensity as an analogue voltage difference through the photoelectric sensor element. Further details regarding the pyrometer settings and data acquisition are given in section 2.3.

2.2. Experimental design

For the experiments, a prismatic geometry with a square cross-section of 10 mm \times 10 mm was sliced using serpentine hatching at a hatch distance of 100 μm , which was rotated by 90 degree per layer. The laser power and scan speed were fixed at 150 W and 800 mm/s throughout, which are default parameters provided by the machine manufacturer. The first 100 layers were processed using the default layer height of 30 μm . As these were intended as sacrificial layers to move the experiment away from the build plate, data from these layers were not included in the training set. Subsequently, different layer heights were applied in sets of 30 layers, running from 10 μm to 110 μm with an increment of 10 μm . It was thought that including 30 layers within each set would mitigate any transitional effects that could arise from the previous layer thickness regions. Therefore, a total of 330 layers were recoated and exposed for inclusion in the data set. Fig. 2 shows the printed specimen's cross-section. One side of the cube was designed with a sawtooth pattern to help identify the start of each new section on the microscopy slides. Dashed line separators show each region, and the layer thickness assignment in each region is given on the right side of the image. The specimen was detached from the build plate using electrode-discharge machining and then bisected in the vertical plane near the centerline using a Struers Accutom-10. The specimens were then embedded, ground with 320-grit sandpaper, and polished with Struers commercial polishing cloths Largo, Dac, Nap, and Chem with suspensions with particle sizes of 9 μm , 3 μm , 1 μm , and 0.1 μm , respectively. Porosity defects, including lack of fusion and other anomalies, were captured using a Keyence VHX-7000 light microscope, which scanned polished mounts under both co-axial and ring lighting




Fig. 1. On-axis pyrometer implementation and the optical path of Aconity Midi + LPBF machine used in the experiment.




Fig. 2. Thickness is changed in every 30 layers, ranging from 10 μm to 110 μm . The nominal value for the parameters used is 30 μm .

modes. Cross-sectional porosity was quantified using ImageJ software by applying a brightness threshold to the captured images. Analysis of the cross-sections confirmed a correlation between powder layer thickness and the severity of porosity-related anomalies, with an increase in layer thickness leading to more pronounced defects.

2.3. Process data acquisition

For this experiment, a dual-channel data collection pipeline was implemented. The primary channel was dedicated to recording the laser trigger signal, which gives info when the laser is on or off, which is crucial for delineating temporal events within the experiment. Simultaneously, the secondary channel was tasked with capturing the pyrometer signal originating from the process zone, which was coaxially aligned and conveyed via a fibre optic cable to the photodiode detector, sourced from Kleiber (model KG740) sensitive in the range of 2 μm to 2.2 μm . The detector has a response time of 6 μs and exhibits a high sensitivity within the wavelength range, which is crucial for accurately capturing emissions from the process zone. The pyrometer is sensitive to the black body emissions arising from the melt pool and the surrounding area, including the powder bed, consolidated metal, and various high-frequency events such as spatter and soot. Although the exact effective diameter around the melt pool for emission measurement is unknown, Planck's law [15], which describes black body emission as proportional to temperature to the fourth power, makes the signal extremely sensitive to the peak temperature. Both channels were sampled at 100 kHz, so the response time of 6 μs was accommodated based on the Nyquist sampling criteria. Central to our methodology was the synchronization of data acquisition across both channels. This synchronization ensured that temporal events, such as the on/off states of the laser trigger corresponding to the scanning strategy employed, were accurately aligned with the emissions detected by the photodiode. Notably, the laser scanning strategy adopted in our experiment resulted in a square wave representation of the laser trigger signal. This facilitated clear temporal correspondence with emissions from the process zone, which were registered as variations in voltage by the photodiode detector. Exemplary pyrometer and trigger signal data are shown in Fig. 3, underscoring the coherence between the two signals. Each square wave of the laser trigger consistently coincided with detectable emissions from the process zone, validating the robustness of our data acquisition setup.

2.4. Pre-processing and preliminary experimental data analysis

Key parts of the dataset where the laser was active were extracted for analysis using the trigger signal and then pre-processed to enhance its quality and relevance. First, the data points were scaled to 0 to 1 to ensure consistency across the entire dataset. Next, the data was divided

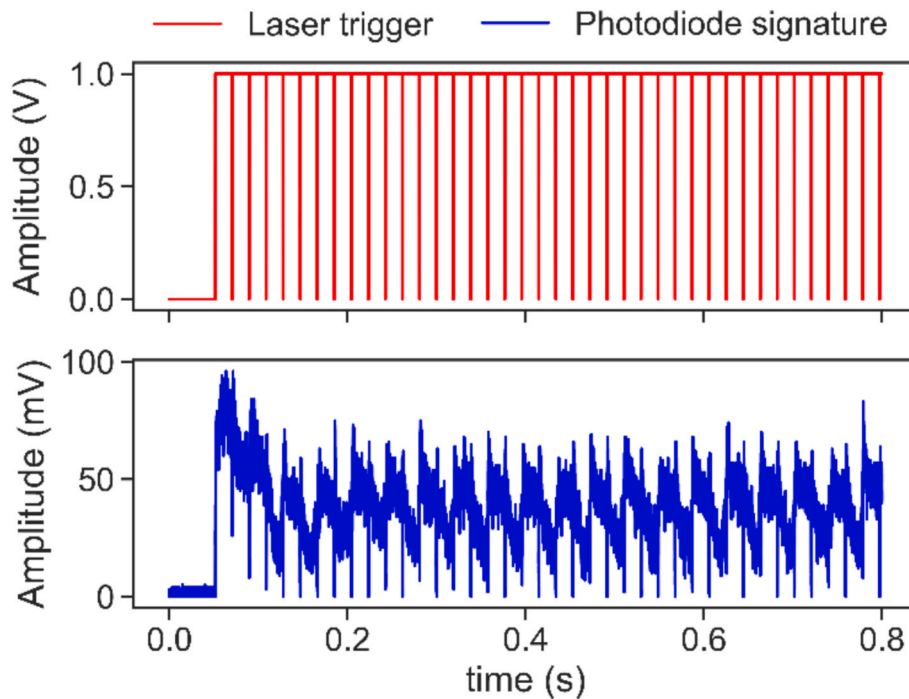


Fig. 3. Laser trigger and optical photodiode signal.

into windows containing 1000 data points. This window size was chosen to balance capturing sufficient temporal detail and maintaining computational efficiency. By segmenting the data this way, the analysis focuses on discrete temporal intervals, facilitating a more detailed investigation of the underlying dynamics. The resulting dataset includes observations across 11 distinct regimes, each representing emissions from the process zone under varying layer thickness. The dataset was balanced with a total of 49,000 windows, with nearly 4400 windows corresponding to each dataset. This comprehensive dataset forms the foundation for statistical analysis, enabling the exploration of process dynamics and potential differences across applied experimental conditions. The analysis primarily focused on the time domain, where windowed signals were utilized to extract relevant time-domain features. This approach enables capturing the temporal variations in waveform characteristics, offering insights into the dynamic behaviour

of the process under investigation, particularly concerning layer thickness variations within the LPBF process. Fig. 4 illustrates the kurtosis distribution as a function of varying layer thickness. Notably, the mean and spread of kurtosis values exhibited variability concerning layer thickness, suggesting distinct temporal dynamics associated with different process conditions.

A closer examination of the kurtosis plot revealed intriguing patterns. Particularly, during instances of very low layer thickness, characterized by pronounced fluctuations in the process dynamics, the distribution of kurtosis values exhibited a wider spread. However, as layer thickness increased, indicative of changes in the process conditions, the spread of kurtosis values tended to decrease. This observation suggests a nuanced relationship between layer thickness and the temporal characteristics of the process dynamics, as captured by the photodiode detector. Overall, the analysis in the time domain provided

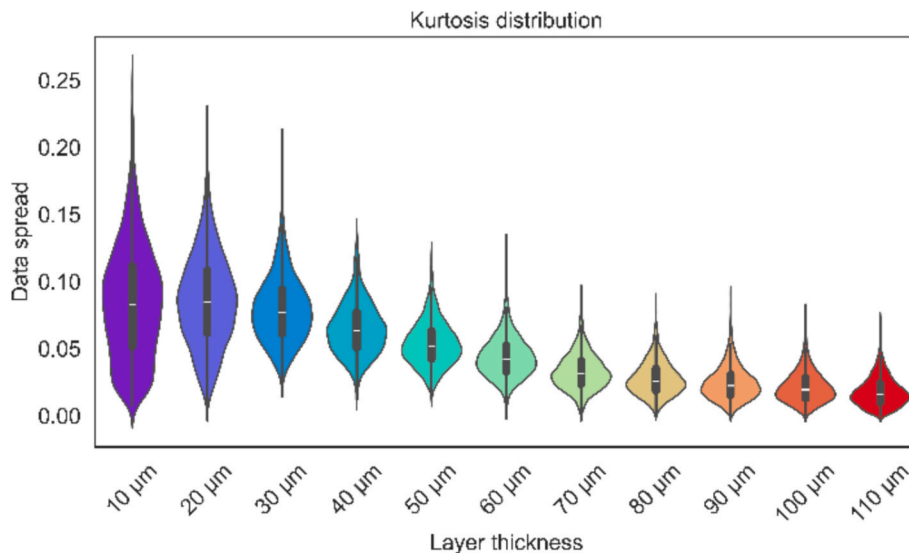


Fig. 4. Kurtosis feature distribution over varying layer thickness with the same laser parameters.

valuable insights into the temporal evolution of the LPBF process, highlighting the influence of layer thickness on the observed dynamics. By systematically examining time-domain features, such as kurtosis, across varying process conditions, a deeper understanding is gained of the intricate interplay between process parameters and the resulting

temporal behaviour, paving the way for informed optimization and control strategies in LPBF applications. Further analysis of another time domain feature, namely skewness (Fig. 5), revealed a trend akin to that observed with kurtosis. Specifically, it was noted that similar to kurtosis behaviour, the skewness values distribution exhibited a wider spread for

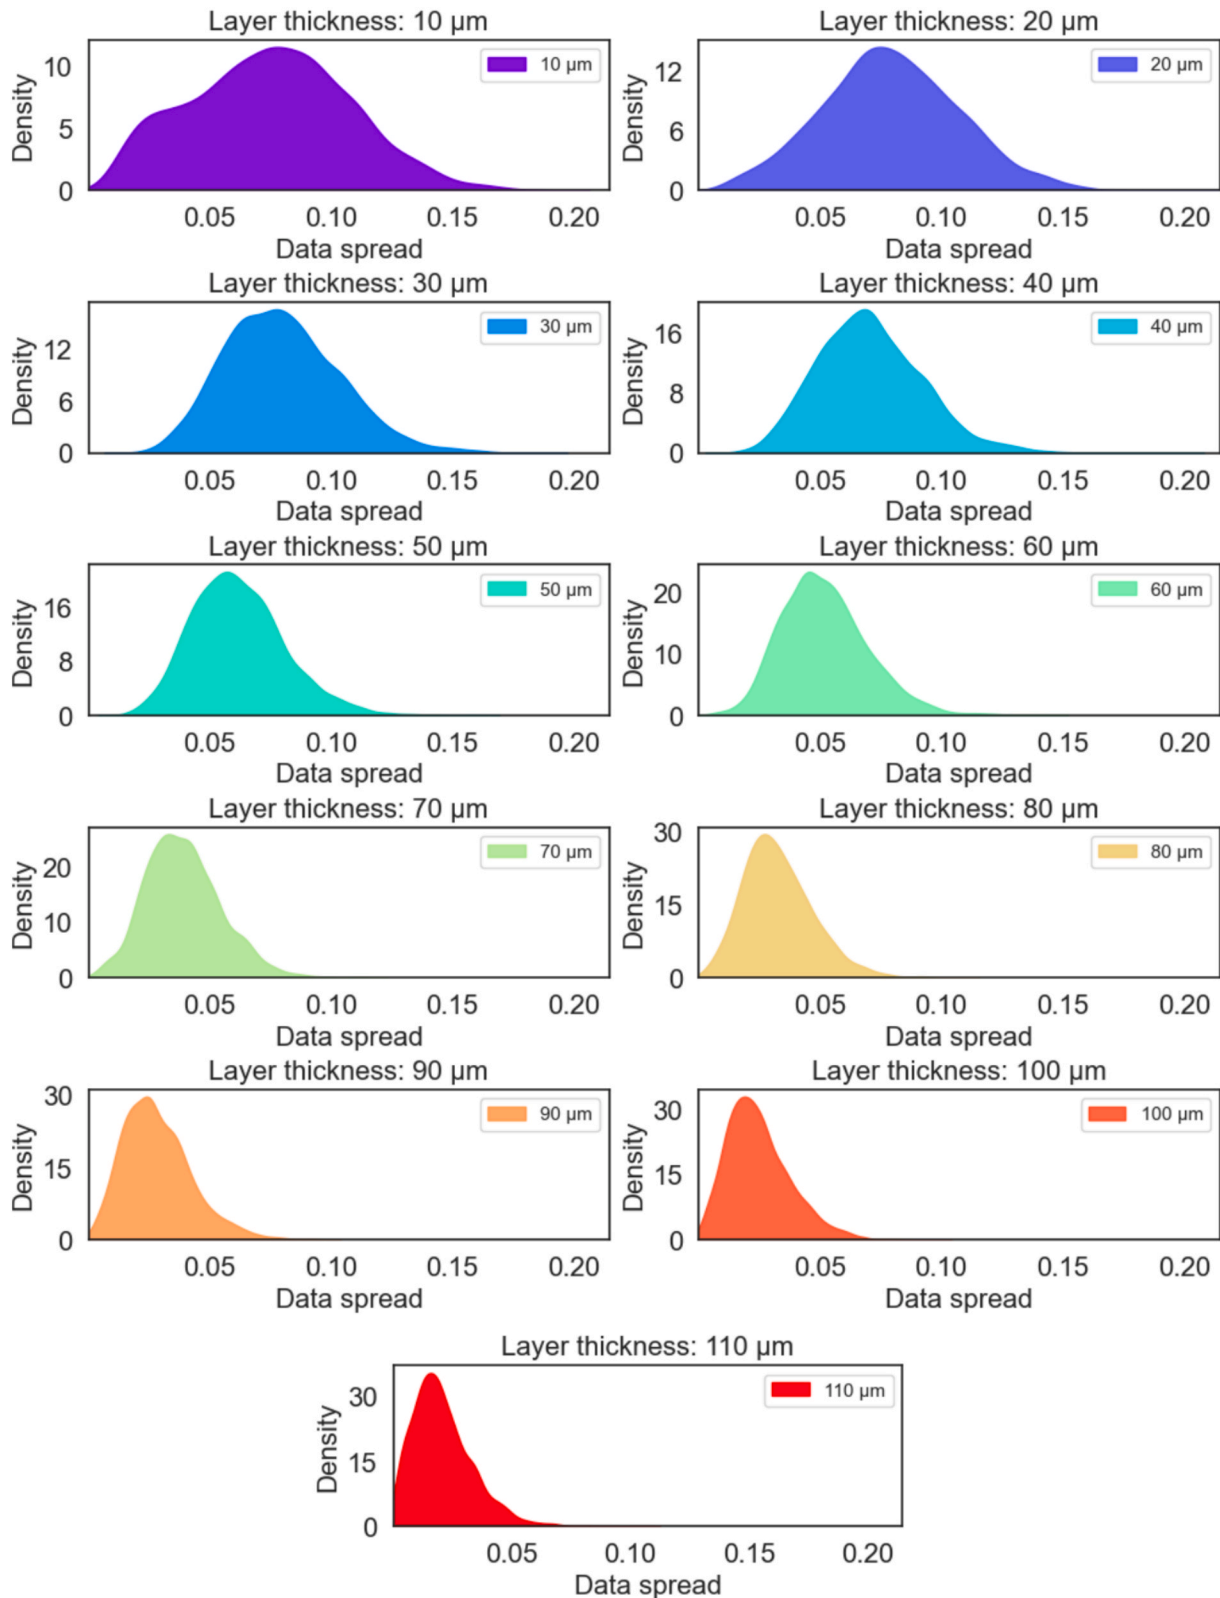


Fig. 5. Skewness feature distribution over varying layer thickness with same laser parameters.

lower layer thicknesses. Conversely, as the layer thickness increased, the distribution became less widespread, accompanied by a noticeable shift in the mean of the distribution.

These observations reinforce that variations in layer thickness affect the temporal signal characteristics captured by the photodiode detector. At lower layer thicknesses, the wider spread of skewness values reflects greater variability in this statistical indicator, whereas at higher thicknesses the narrower distribution corresponds to a more constrained range. A closer examination of time-domain statistical features thus provides clear evidence of the photodiode detector’s ability to capture the nuanced dynamics of the LPBF process in response to changes in layer thickness. We gain deeper insights into the complex interplay between process parameters and temporal behaviour by systematically analyzing time-domain features such as skewness in conjunction with kurtosis and other relevant time domain metrics. Like the analysis conducted in the time domain, frequency domain features were analyzed to further elucidate the dynamics of the process, particularly in response to varying layer thicknesses. The sampling rate of 100 kHz results in a Nyquist frequency of 50 kHz, thus ensuring that the frequency components within the signal were adequately sampled. Frequency domain analysis identifies repetitive patterns within the waveform and quantifies their energy across different frequency bands. This analysis aimed to ascertain whether there were any discernible patterns or changes in energy distribution across frequencies as a function of layer thickness. To achieve this, the Welch periodogram method was employed, where the signal was partitioned into five frequency bins: 0 kHz to 10 kHz, 10 kHz to 20 kHz, 20 kHz to 30 kHz, 30 kHz to 40 kHz and 40 kHz to 50 kHz. The energy of the waveform within each frequency window was computed across the entire dataset and plotted against layer thickness. The resulting distribution of energy components across frequency windows revealed notable insights as shown in Fig. 6. Most of the frequency components primarily resided within the lowest frequency range, with the highest energy concentration observed in the band ranging from 0 kHz to 10 kHz. Furthermore, upon closer examination, it became evident that the energy of low-frequency components increased with layer thickness.

This trend of increasing energy with layer thickness was consistent across all frequency windows, albeit less pronounced in higher frequency bands. Table 1 summarizes the results of one-way ANOVA tests performed on selected time-domain and frequency-domain features

Table 1

One-way ANOVA results for key time-domain and frequency-domain features extracted from signals across 11-layer thickness classes (10 μm to 110 μm).

Feature	Time or Frequency domain	F-statistics	P-value	Statistically significant
Kurtosis	Time	141.03	0.00245	Yes ($p < 0.05$)
Skewness	Time	1179.53	0.00009	Yes ($p < 0.05$)
Frequency energy (0–10 kHz)	Frequency	623.65	0.00013	Yes ($p < 0.05$)

extracted. The analysis was conducted across 11 distinct layer thickness classes (10 μm to 110 μm) to assess whether the features exhibited statistically significant variation with respect to the class. The F-statistic values represent the ratio of between-class to within-class variance, where higher values indicate more substantial evidence of class-dependent differences. All three features—kurtosis, skewness, and frequency energy in the 0–10 kHz band—demonstrated statistically significant differences across the classes ($p < 0.05$), confirming that these features effectively capture thermal behaviour of the melt pool due to changes associated with varying layer thicknesses.

The comprehensive analysis conducted across both time and frequency domains both with visualization plots and one-way ANOVA serves as a robust indicator of the photodiode detector’s ability to detect thermal fluctuations and perturbations within the process zone. This holistic understanding not only underscores the efficacy of the photodiode detector for monitoring process dynamics but also informs strategies for optimizing and facilitating the identification and stabilization of critical process parameters in LPBF applications to achieve desired outcomes.

3. Machine learning framework

The preliminary analysis demonstrated statistically significant variations in process emissions that are strongly correlated with layer thickness, making these emissions valuable indicators for real-time monitoring of LPBF build quality. As previous studies have done, these emissions provide critical feedback that can be used to detect and prevent defects such as keyhole porosity and lack-of-fusion during the build process [12]. As reported in the literature, traditional monitoring

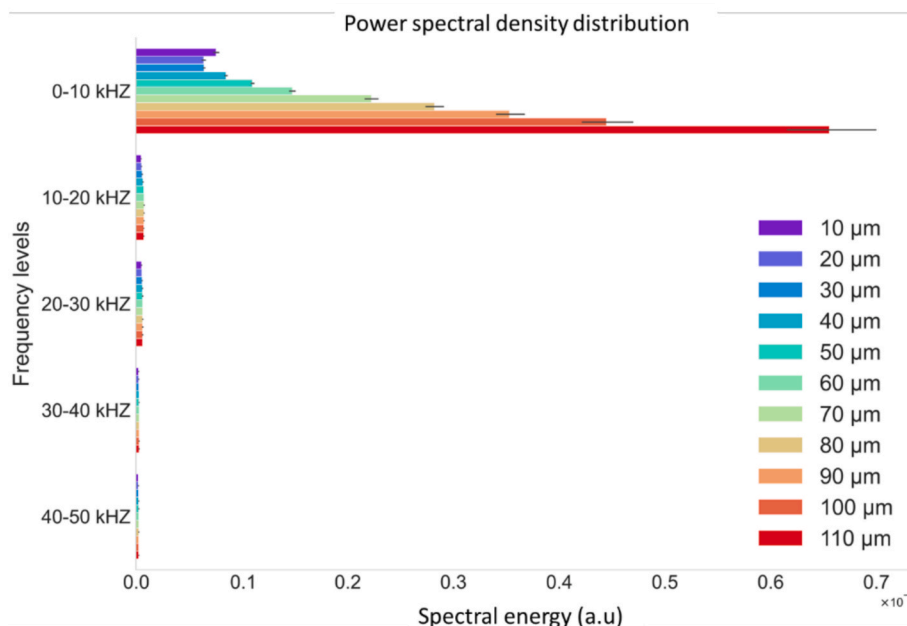


Fig. 6. Power spectral density distribution over varying layer thickness on five different ranges (values are not normalized).

methods often rely on machine learning models trained on signals of fixed length. These models are built to process inputs of a specific size or time scale. As a result, each time the scan length or vector geometry changes, a separate model must be trained for that particular configuration. This creates a cumbersome and disjointed workflow, requiring developing and managing multiple models for different scanning strategies and geometries. Such fixed-length signal-based models struggle to adapt to the inherent variability in laser scanning strategies, where scan lengths and vector patterns can differ significantly depending on the

part geometry. This makes the monitoring process less efficient and limits the broader applicability of these models in real-world LPBF applications, where dynamic adjustments to scanning parameters are common. To address this limitation, we introduce a neural network model that overcomes the need for fixed-length inputs. Rather than being constrained by the scan length or vector geometry, our model is designed to focus on the emission characteristics directly correlated with layer thickness, which are the most critical factors for build quality. The model is inherently insensitive to variations in scanning patterns,

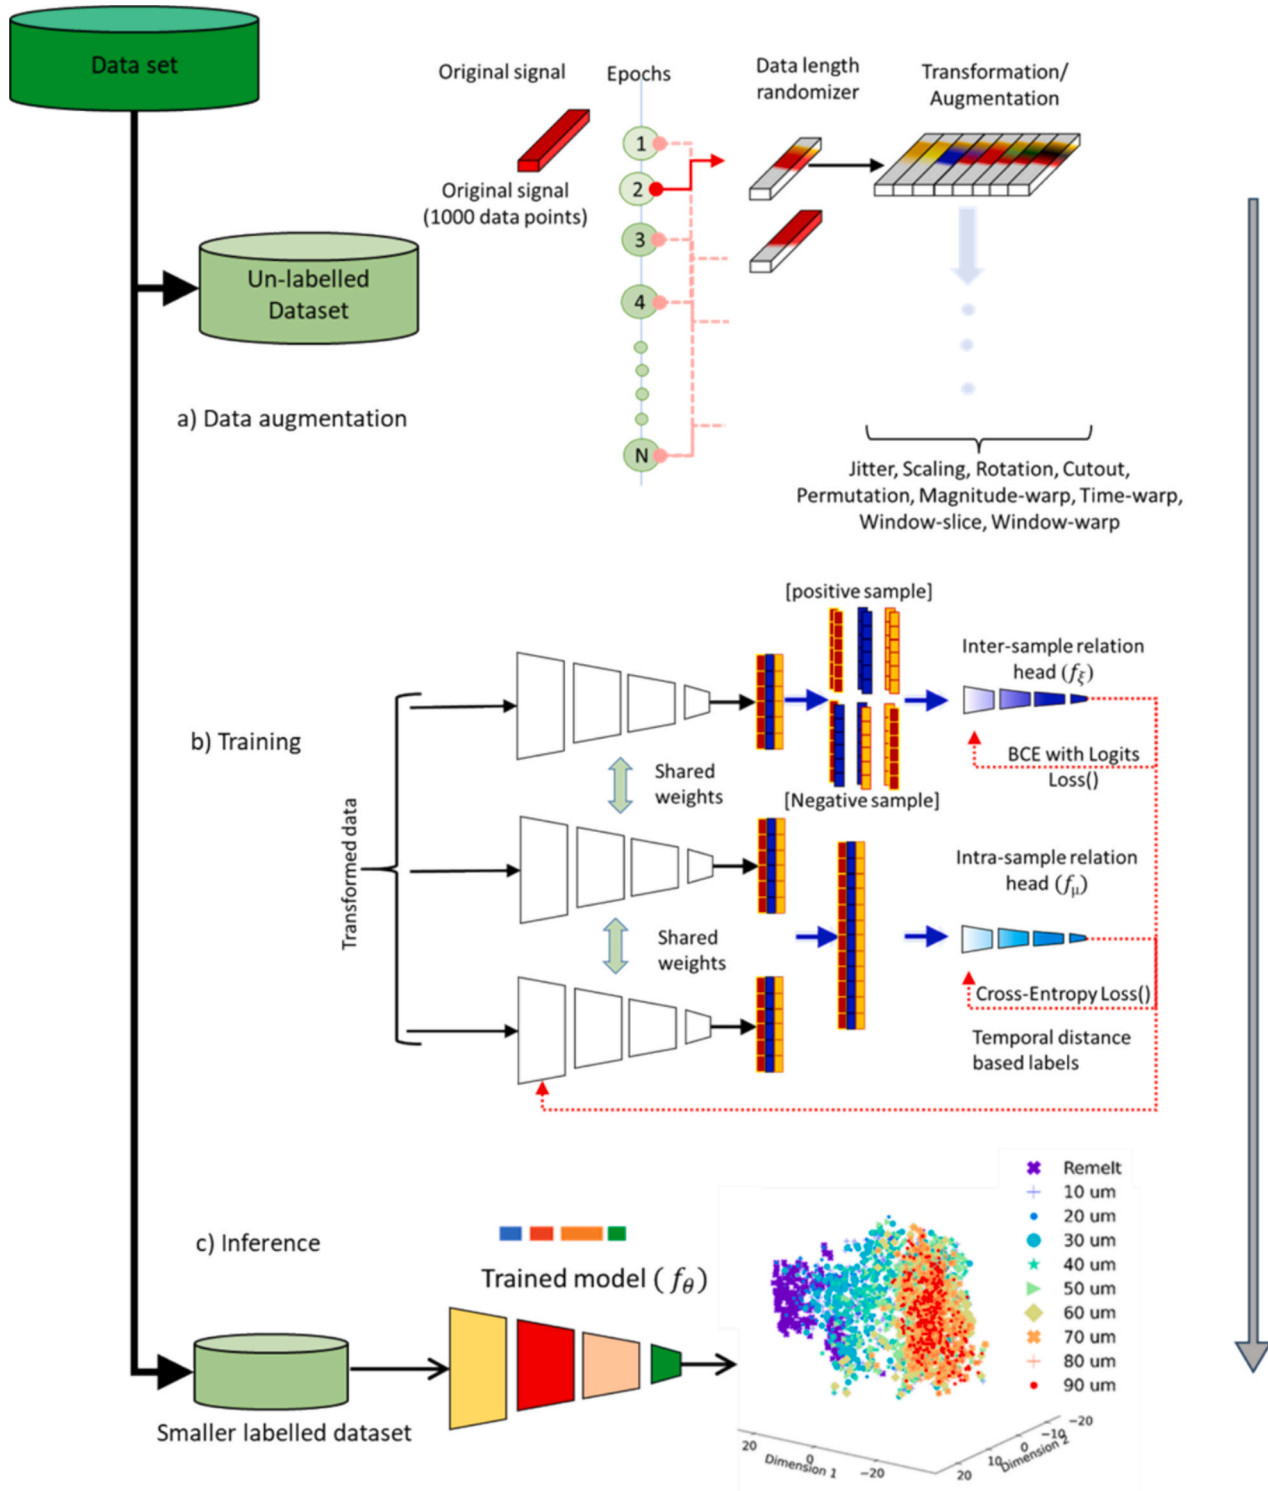


Fig. 7. Schematic of the proposed methodology.

enabling it to generalize across different part geometries and laser scan strategies without needing retraining for each configuration. As an added feature, we train this model in a self-supervised manner, which provides an additional advantage by addressing one of the significant challenges in machine learning for manufacturing: data labelling. Acquiring high-quality labelled datasets is costly and time-consuming in many real-world manufacturing scenarios. Supervised learning approaches require extensive labelled datasets, which are often difficult to obtain, especially when capturing subtle defects or irregularities. Retraining or updating the model may also require additional labelled datasets, limiting this approach's general scalability and applicability. The reliance on manual labelling slows the model development process and introduces human error and bias.

Our model learns to identify patterns and anomalies in process emissions without needing labelled data by adopting an unsupervised learning framework. This dramatically simplifies the implementation of monitoring systems, as it bypasses the labour-intensive task of data labelling. The model can autonomously detect variations and deviations in the emissions that correlate with quality issues, offering a scalable and adaptable solution for real-time process control in LPBF. The combination of self-supervised learning and the model's generalization across different scan parameters ensures that our approach is both efficient and robust, making it a versatile tool for improving LPBF monitoring systems. The pipeline for training this model comprises three stages: Data pre-processing, model training, and inference as depicted in Fig. 7. Data is randomly assigned to two subsets: seventy percent for training the model without access to ground truth information and thirty percent containing ground truth data. Despite possessing prior knowledge of ground truths associated with the dataset, the research methodology adopts a novel approach to analysis. A temporal encoder (f_θ) is trained in a self-supervised manner, marking a departure from conventional methods and introducing a paradigm shift in data processing and interpretation. What sets this encoder (f_θ) apart is its ability to effectively handle the variable time scales inherent in optical waveforms. Through self-supervised training, the encoder (f_θ) learns to extract meaningful patterns and features from the optical signals without relying on ground truth information. Another novelty of the work lies in the temporal network architecture of the encoder (f_θ), which enables it to generate generalized representations of input waveforms independent of their length. After obtaining sparse representations, a critical evaluation of their efficacy is conducted through visual inspection and plotting, providing an initial assessment of the encoder's performance in capturing intricate dataset nuances.

A lower-dimensional embedding technique, such as t-distributed stochastic neighbour embedding (t-SNE), is employed to elucidate further the relationship between sparse representations and underlying powder layer characteristics. This technique helps reveal the inherent structure and relationships within the data with straightforward interpretation. Data corresponding to extreme, intermediate, and minimal powder layer thicknesses in the lower-dimensional space, alongside their respective ground truth labels, are then extracted. The computation of centroids for these distributions provides crucial reference points for subsequent analysis. By leveraging this embedding technique, the proximity of new signals to computed centroids is assessed, facilitating signal classification into distinct categories based on predefined thresholds. The classification process not only distinguishes between powder layers of varying thicknesses but also provides insights into phenomena such as remelting. By utilizing the learned representations and spatial relationships encoded within the t-SNE embedding, informed decisions regarding powder layer depth and characteristics can be made, thereby advancing understanding of the complex dynamics inherent in powder bed additive manufacturing processes. The neural network architecture employed in this study comprises an encoder (f_θ) and two heads, dedicated to modelling inter-sample (f_ξ) and intra-sample (f_μ) relations, respectively as given in Fig. 7.

The proposed approach derives representations on unlabelled

waveform vectors by exploring relationships between different vectors and within the same vector's temporal segments. This process involves the use of a shared encoder (f_θ) backbone for representation learning, based on a TCN. The challenge of data sparsity in additive manufacturing is commonly addressed through techniques such as synthetic defect generation using generative models, or physics-informed modelling that embeds prior domain knowledge into learning frameworks. In this work, we adopt a complementary approach by employing temporal data augmentation strategies that manipulate waveform segments through scaling, jittering, and slicing operations. These augmentations increase the diversity of the training samples without requiring additional labelled data, thereby improving the model's ability to generalize across different process regimes. Importantly, this augmentation operates entirely in the self-supervised domain, aligning with our broader objective of reducing reliance on annotated defect datasets. By enriching the signal variability, the framework becomes more resilient to unseen process conditions and is better suited for real-world deployment where defect-labelled data is scarce or imbalanced. For inter-sample (f_ξ) relation reasoning, a particular sample t_i is chosen as a reference point and subjected to standard augmentation techniques $A(t_i)$ to generate a corresponding positive sample pair (t_i, t_{k-1}) by employing the time series data augmentation operations as depicted in Fig. 8, where k is the number of augmentation. In this case, seven augmentation operations were applied: jitter (adding noise), scaling (adjusting amplitude), cutout (randomly dropping parts of the signal), magnitude-warp (modifying the intensity of certain signal regions), time-warp (distorting the timing of signal events), window-slice (extracting smaller segments from the signal), window-warp (distorting a segment within a window) and combination of all augmentations.

Conversely, another sample t_j from a different time scale, augmented with the same function $A(t_j)$, serves as a negative sample pair (t_j, t_{k-1}) . The relationship reasoning between these samples is conducted on the encoded representations. $z_i = f_\theta(t_{k-1})$ for positive and $z_j = f_\theta(t_{k-1})$ for negative samples by integrating them with the anchor sample's representation $z_i = f_\theta(t_i)$ through a relation reasoning mechanism, which is also a CNN encoder head (f_ξ). The inter-sample (f_ξ) relation between each vector is thereby framed as a binary classification task, with the outcome determined by the output of the relation reasoning head, $h_{(i,j)} = f_\xi(z_i \cdot z_j)$, and trained using a binary cross-entropy loss function L_{inter} :

$$L_{inter} = - \sum_{n=1}^{2N} \sum_{i=1}^k \sum_{j=1}^k (y_n^{(i,j)} \cdot \log(h_n^{(i,j)}) + (1 - y_n^{(i,j)}) \cdot \log(1 - h_n^{(i,j)})), \quad (1)$$

Here, $y_n^{(i,j)}$ is equal to 0 for positive pairs and 1 for negative pairs. To capture intra-temporal relationships within a single vector, $t_i = (t_{i,1}, t_{i,2}, \dots, t_{i,L})$, the sample is divided into equal segments with overlaps. Segments are then categorized as anchor, short-term, mid-term, or long-term based on their temporal proximity to the anchor segment.

By selecting two adjacent subsequences $S_{i,u}$ and $S_{i,v}$ from t_i , their relationship C is defined through the temporal gap $d_{u,v} = |u - v|$, with a distance threshold D set dynamically as $\lfloor L(\text{samplelength})/C \rfloor$. The classification of relationships is based on the ability to fit D within u, v . Intra-temporal reasoning then processes the encoded representations $z_{i,u} = f_\theta(S_{i,u})$ and $z_{i,v} = f_\theta(S_{i,v})$ integrating them through a CNN intra-relation head (f_μ), culminating in a multi-class classification as below:

$$L_{intra} = - \sum_{j=1}^K \left(y_{(u,v)}^j \cdot \log \frac{\exp(h_i^{(u,v)})}{\sum_{c=1}^C \exp(h_i^{(u,v)})} \right), \quad (2)$$

In this equation, $y_{(u,v)}^j$ represents the temporal labels. This approach enhances the understanding of patterns within time series data, enabling more profound insights into the interrelations among various data points and temporal segments, significantly benefiting high-dimensional time series data analysis. The combined model training utilizes an

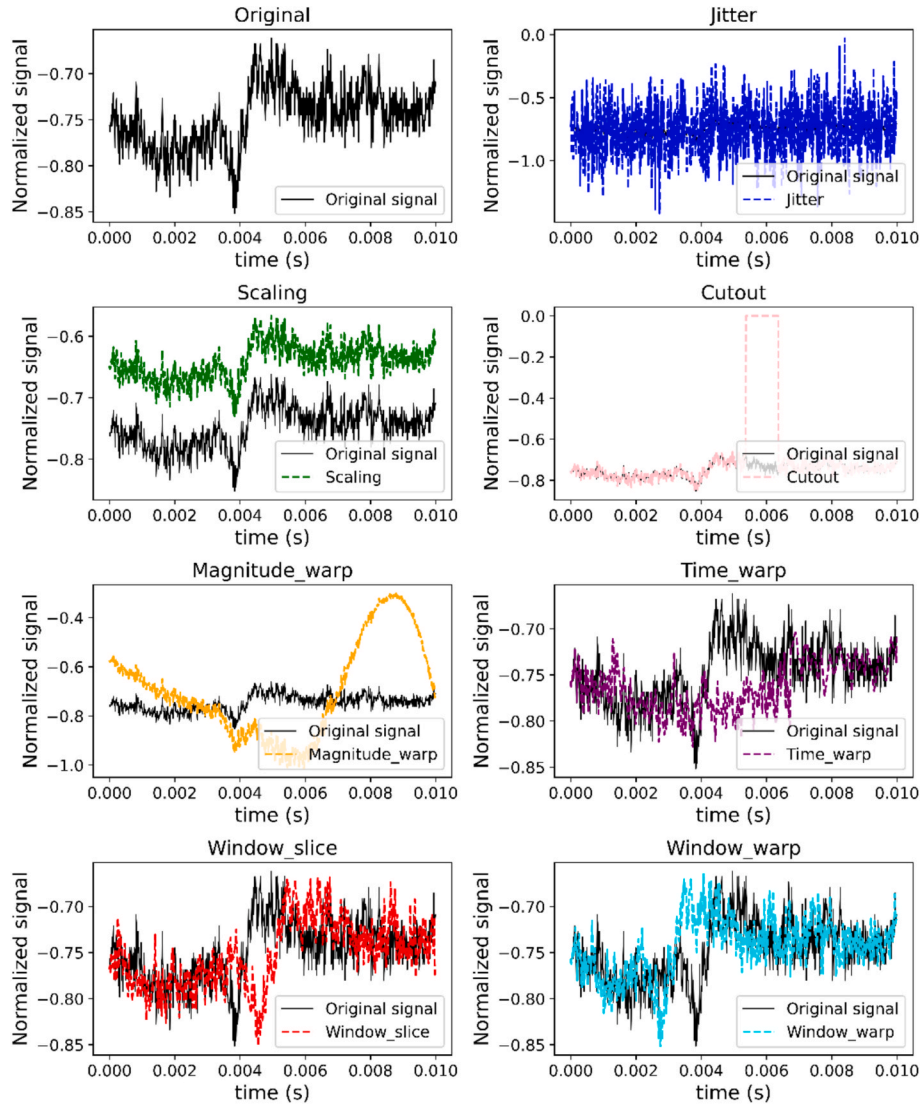


Fig. 8. Augmentation applied on original waveforms.

overall loss function, facilitating a comprehensive learning process.

$$L_{\text{overall}} = L_{\text{inter}} + L_{\text{intra}} \quad (3)$$

3.1. Temporal convolution network (TCN) architecture

TCN is a crucial architecture in learning representations from sequential or time-varying data, particularly in tasks like time series analysis. TCNs efficiently capture local patterns and long-term dependencies using dilated convolutions and hierarchical feature extraction, enabled by multiple layers and residual connections. By leveraging parallelization and temporal pooling, TCNs effectively condense information, making them adaptable to various tasks beyond time series analysis. Their inherent flexibility, scalability, and ability to capture temporal dependencies make TCNs essential for achieving state-of-the-art performance in sequential data modelling tasks. The encoder backbone architecture is based on TCNs and underwent optimization through an exhaustive search procedure. The architecture is shown in Fig. 9. This backbone comprises three blocks within the TCN architecture, each consisting of two 1-dimensional convolution layers. In the first TCN block, the input tensor has a channel size 1 and an output channel size 8. The data size is restored after a kernel operation of 16 by using dynamic padding and a dilation size of 1. The second TCN block is similar to the first, with a dilation size of 2, taking the output from the

first TCN block. The third TCN block is akin to the first two, with a difference in the output channel size of 64. It also receives input from the first and second TCN blocks, acting as a residual connection. Subsequently, adaptive pooling of 1 was applied to the output of the third TCN block to produce a single-dimensional output across all channels. In this architecture, the initial temporal convolution layer receives a tensor with dimensions $B \times 1 \times L$ (where B represents the batch size and L denotes the length of the optical signal), generating a feature map sized $B \times 64$. This feature map serves as the sparse representation of the optical waveform signal. Each temporal convolutional layer employs a filter size of 16. With approximately 90,000 trainable parameters, the backbone model was implemented using the PyTorch framework, utilizing PyTorch's built-in activations and max-pooling operations for ReLU activations and max-pooling operations, respectively.

3.2. Inter-sample and intra-temporal relation heads

As a result, the TCN encoder backbone (f_{θ}) generates representations for each triplet containing an anchor, positive, and negative sample, enabling subsequent analysis and processing. Additionally, randomly augmenting each sample 8 times increases the batch size 8 times. Constituting integral components of the methodology, the inter-sample head (f_{ξ}) and intra-sample (f_{μ}) head consist of two linear, fully con-

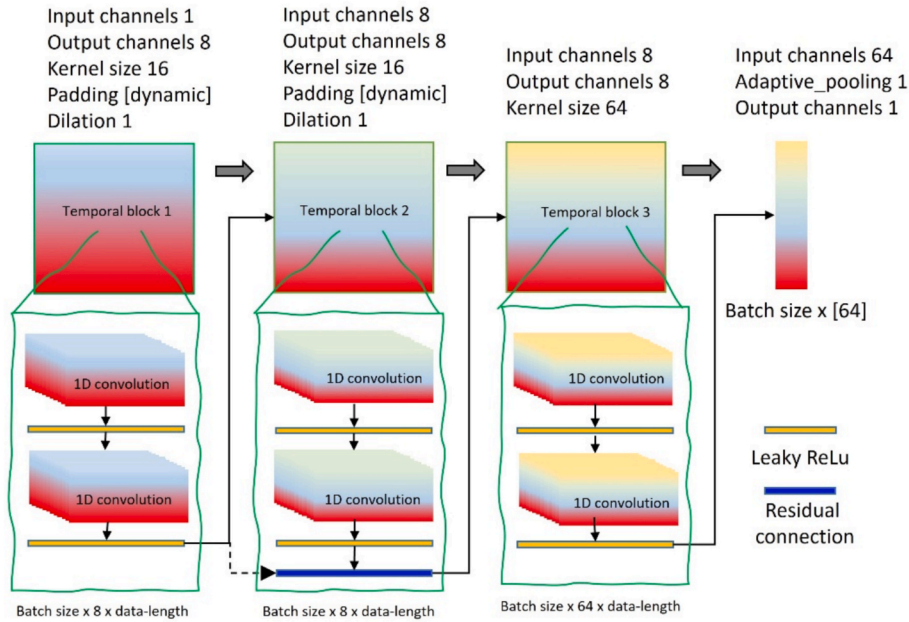


Fig. 9. TCN architecture (f_θ) proposed in this work.

nected layers. These elements play a crucial role in harnessing the TCN encoder’s latent dimension of 64, initiating a comprehensive exploration into representation learning. Initially featuring congruent architectural structures, both heads include a foundational layer comprising 256 neurons, as illustrated in Fig. 10. However, a subtle distinction arises in the second layer: while the inter-sample head f_ξ produces a single output, the intra-sample head (f_μ) is responsible for computing outputs that capture temporal relationships, specifically calibrated to a temporal span of three within our study.

The inter-sample (f_ξ) head performs a critical function in executing binary classification assessments orchestrated by an inter-sample (f_ξ) relation linear neural network, culminating in the quantification of misclassification through the loss function (L_{inter}). Subsequently, the narrative delves into the realm of intra-temporal relation reasoning, a complex endeavour involving the segmentation of each sample into equidistant segments. These segments undergo dynamic categorization based on their temporal proximity to an anchor segment. This nuanced approach assigns short-term, middle-term, or long-term relationship designations to each segment, thereby effectively shaping the narrative of temporal dependency. The next stage entails reformulating the intra-sample (f_μ) relation reasoning task into a multi-class classification problem using an intra-sample (f_μ) relation neural network, leveraging representations derived from the encoder backbone. At this pivotal juncture, the loss function (L_{intra}) is computed to evaluate misclassification. Subsequently, through back-propagation and weight adjustment,

the combined loss given by equation (3) is propagated backwards to update the weights of the encoder and the two heads. This iterative process continues until the encoder acquires discernible representations within the data, thereby facilitating effective representation learning for sequential or time-varying data.

3.3. Data randomizer

The core novelty of this research lies in the innovative strategy designed to handle waveforms with variable data lengths, which is essential for effectively training TCN, as depicted in Fig. 11. A dynamic data loader was developed to address this challenge, enabling flexible data handling within the TCN encoder backbone f_θ framework. The dataloader randomly selects a segment length between 500 and 999 data points from an original time series signal consisting of 1000 points. This random selection ensures that the model does not become dependent on a fixed window size, allowing it to generalize better across varying signal lengths and dynamic conditions. Once the segment is selected, the dataloader truncates the waveform to the specified length. This truncation step mimics real-world scenarios where time series data might be incomplete or noisy, ensuring that the TCN network is robust to such variabilities. Importantly, this dynamic truncation helps simulate the random nature of real-world signals and improves the model’s ability to handle diverse process variations. After truncation, the data undergoes various augmentations, as described above, within the dynamic data

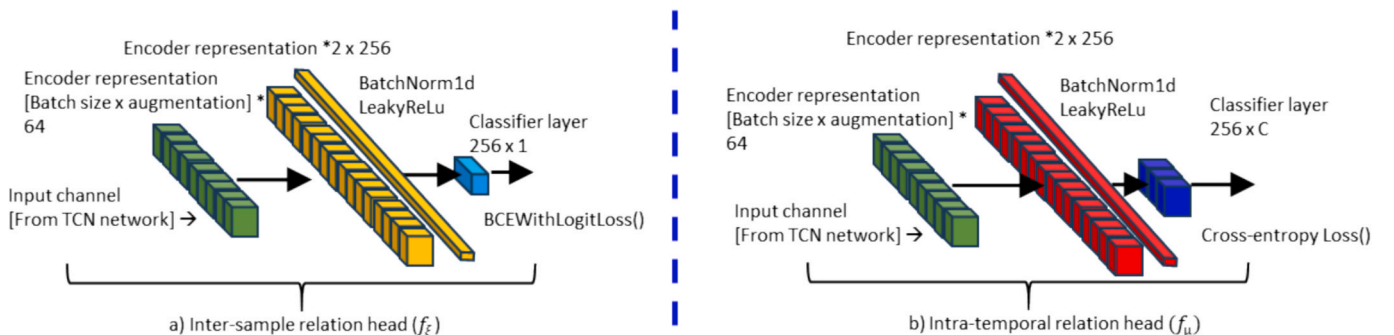


Fig. 10. Heads to evaluate the temporal relationship.

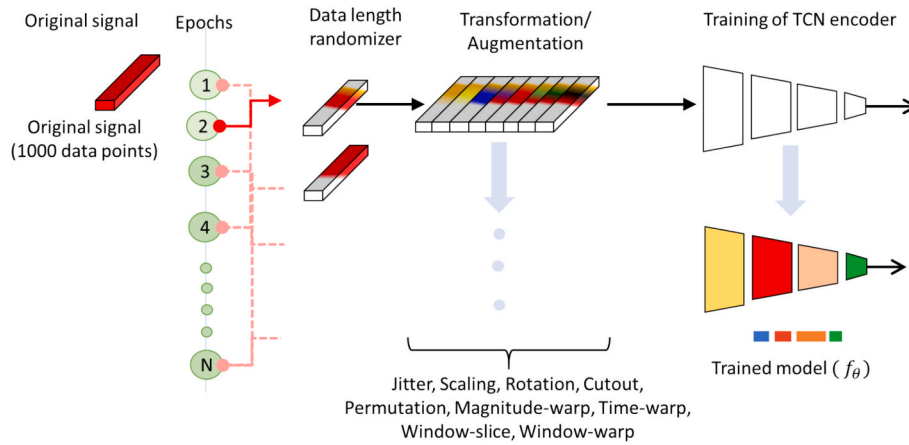


Fig. 11. Data randomizer across epochs.

loader. These transformations enrich the dataset by creating new variations of the original waveforms, effectively augmenting the data available for training. The augmented and truncated data is then passed through the dynamic data loader into the TCN encoder backbone f_θ network, ensuring the model is exposed to a wide range of waveform distortions and sizes during training. This augmentation process not only improves the diversity of the dataset but also enhances the TCN network's ability to learn robust representations that generalize well across different tasks.

4. Results

The following section discusses the training curves, performance evaluation, and the effectiveness of the encoder in detail after training. The TCN encoder backbone f_θ trained in a self-supervised manner would learn the representation of the waveform signal. Subsequently, this backbone can be employed as a feature extractor to carry out downstream tasks, like anomaly detection, classification or regression for the observed variable.

4.1. Model training

The three models, consisting of the encoder (f_θ) and the two heads (f_ξ , f_μ), underwent concurrent training using the Adam optimizer, a widely utilized optimization algorithm in deep learning. The Adam optimizer adjusts the model parameters based on the gradients of the loss function to those parameters, enabling efficient optimization. During training, the learning rate, which determines the step size of parameter updates, followed a cosine annealing schedule. This gradually reduced the learning rate from an initial value of 0.01 to a final value of 0.001 throughout training. This technique helps to fine-tune the model more effectively and potentially overcome local minima in the optimization landscape. A batch size of 256 was utilized during training, indicating the number of samples processed in each iteration of the optimization process. Larger batch sizes can lead to more stable training and faster convergence but require more memory and computational resources. Training of the models occurred on a powerful hardware-accelerated Graphical Processing Unit (GPU), specifically a NVidia RTX Titan. GPUs are well-suited for training deep neural networks due to their parallel processing capabilities, significantly speeding up computation compared to traditional CPUs. The three models were trained for a total of 200 epochs, where each epoch represents one complete pass through the entire training dataset. This prolonged training duration allows the models to learn from the data and adjust their parameters to minimize the cumulative binary (L_{inter}) and multi-label (L_{intra}) cross-entropy loss values. The training parameters used for updating the weights of the backbone and the two relation-reasoning heads were carefully chosen to

balance model convergence and generalization (Table 2). These parameters include the learning rate schedule, batch size, and optimization algorithm settings, all of which play critical roles in determining the effectiveness and efficiency of the training process.

During training, the unlabeled optical waveform of various time scales was taken from the dataset and split into three subsets: 60 % for self-supervised training of the TCN encoder backbone, 10 % for plotting the performance curves and 30 % for supervised anomaly centroid based algorithm training and model robustness evaluation. The distribution of variable windows utilized for training visually demonstrates that the windows used to train the model were spread evenly across the selected ranges, indicating a lack of bias in the training data (Fig. 12). This distribution plot provides a clear overview of the temporal scales present in the optical windows and highlights the diversity of data used to train the model. By ensuring that the training data covers a wide range of temporal scales without bias towards specific ranges, the model is better equipped to learn robust representations and generalize effectively to unseen data. This balanced distribution of variable windows contributes to the overall reliability and robustness of the trained model.

The graphical representation in Fig. 13 illustrates the progression of cumulative loss (L) and training accuracy for 200 epochs. It is evident from the plots that the cumulative loss magnitude and accuracy stabilize around the 150th epoch, suggesting diminishing performance improvements beyond this stage. This observation indicates that the temporal encoder backbone model effectively learns the distributions of optical signatures from the data.

The stabilization of the cumulative loss magnitude signifies that the model has reached a point of convergence where further training epochs do not yield significant performance improvements. This finding underscores the effectiveness of the training process in capturing and encoding the essential features of the optical data. Overall, the results suggest that the Temporal encoder backbone model has successfully

Table 2

Training parameters of the models trained with self-supervised loss.

Training parameters	Parameters
Type of analysis	Representation learning
Solver name	'Adam'
Learning rate	0.01 – starting, 0.0001 – stopping
Batch size	256
Data transformation	True
Number of augmentations	8
Number of temporal relations C	3
Shuffle	Every epoch
Training set	60 % of the dataset (Unlabelled)
Validation set	10 % of the dataset (labelled)
Training epochs	200
Trainable weights	$\approx 90\,000$




Fig. 12. Learning rate and variable window distribution used for Model training.




Fig. 13. Learning Curves.

learned meaningful representations of the optical signature distributions, highlighting its capability to extract valuable information from the input data. This finding is crucial for ensuring the reliability and efficacy of the model in subsequent tasks, such as classification or clustering of optical waveforms.

4.2. Model inference with representation

Another validation method for assessing the effectiveness of the trained temporal encoder backbone model involves visualizing its lower-dimensional representation alongside known ground truth data. Although the ground truth was not used during training, it serves as a reference for inference. Any discrepancies between the model’s representation and the ground truth distribution indicate unsuccessful learning of temporal relationships within the data. This qualitative assessment compares the model’s learned representations and the ground truth, providing insights into whether the model has effectively captured the underlying temporal dynamics. Fig. 14 shows the distribution plots across all 32 latent spaces, revealing non-overlapping statistical differences in mean distribution, which indicates a significant shift across individual latent vectors. This illustrates the model’s proficiency in capturing underlying patterns and features within the data. Demonstrating distinct statistical differences across latent vectors underscores the model’s ability to discern and encode relevant temporal information, further validating its efficacy in representation learning. Furthermore, to deepen the understanding of the learned representations across variable lengths, *t*-SNE is employed. This technique decomposes the 32-dimensional latent representations into lower-dimensional spaces across four labelled datasets with four different window lengths: 1000, 900, 750, and 500 data points. The resulting 2D visualizations in Fig. 15 comprehensively depict the distribution spread

across the data spaces, revealing clusters corresponding to the 11 distinct process conditions (i.e., layer thicknesses) in the dataset. Additional analysis of the clustered space on the training set’s optical signatures over the layer thickness further validates the model’s efficacy. Moreover, overlaps between adjacent layer thicknesses are observed, indicating transitional dynamics within the underlying process. This observation underscores the nuanced complexities inherent in the data. The TCN encoder implicitly encodes the thermophysical behavior of the melt pool by generating feature embeddings that exhibit a smooth and ordered distribution aligned with layer thickness variations, as revealed through *t*-SNE visualization—despite not being trained on explicit physical models or governing equations. The *t*-SNE manifold illustrates that the learned representations span a continuous space across layer thicknesses, indicating that the process space is continuous and that the framework effectively probes a spectrum of thickness variations rather than treating the problem as a strictly discrete classification task. Overall, the results highlight the capability of the temporal encoder backbone to capture and represent the complex structures inherent in sequential data. Building upon the success of distance-based metrics, the analysis was extended to compare centroid distances across all classes. This approach allowed the exploration of the evolution and spread of the distance metric between classes over different window lengths.

The first step in the analysis involved calculating the centroid for each class. In this case, the centroid is a single point in the three-dimensional feature space representing the mean position of all data points (samples) in the class. Mathematically, the centroid of a class C_i is given by.

$$C_i = \frac{1}{n_i} \sum_{j=1}^{n_i} x_j \quad (4)$$

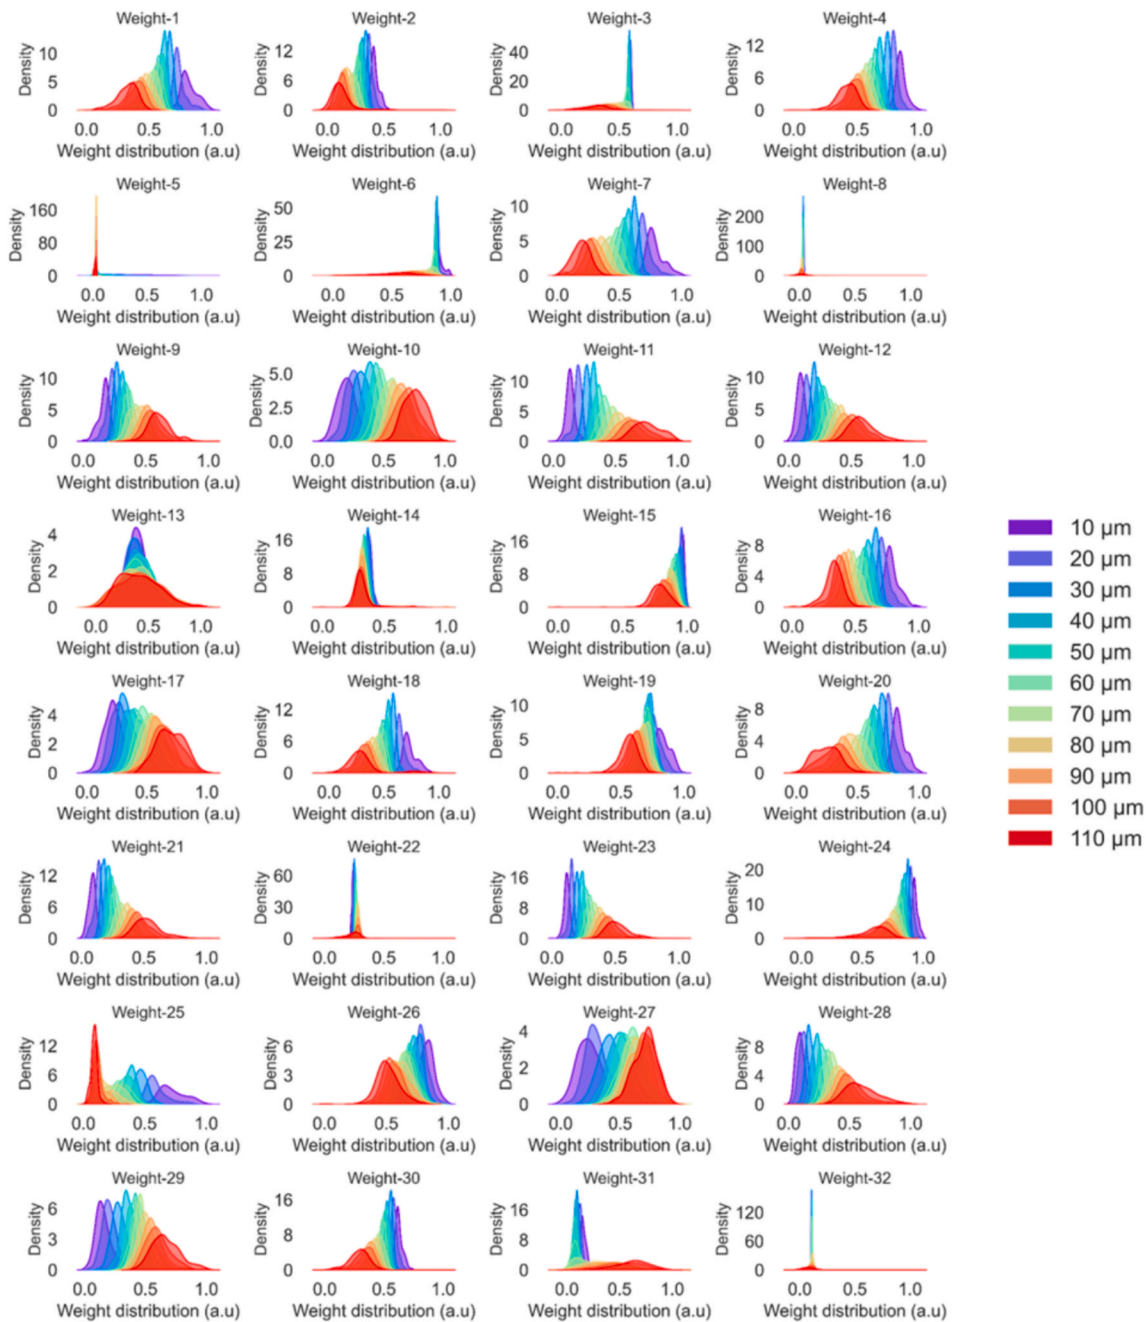


Fig. 14. Distribution of latent dimension across the different layer thicknesses.

where n_i is the number of samples in class C_i , and x_j is the feature vector of the j^{th} sample within the class. This step is critical because the centroid serves as a reference point that summarizes the overall characteristics of the class. By computing the centroids for all 11 classes, we can reduce the complexity of analyzing individual samples and instead focus on comparing the average behaviour of each class. Once the centroids were computed, the Euclidean distance between the centroid of each class and every individual data sample was calculated. This includes the samples within the same class and those in other classes. The Euclidean distance between a class centroid C_i and an individual sample x_j is defined as:

$$d(C_i, x_j) = \sqrt{(x_{j1} - C_{i1})^2 + (x_{j2} - C_{i2})^2 + (x_{j3} - C_{i3})^2}, \quad (5)$$

where C_{i1}, C_{i2}, C_{i3} are the coordinates of the centroid, and x_{j1}, x_{j2}, x_{j3} are the corresponding feature values of the sample. This step provides a way

to measure how far each data point is from its respective class centroid and how far it is from the centroids of other classes.

By calculating these distances, one can identify samples that are close to or far from their centroid, giving insights into potential outliers or samples that are more representative of their class. In addition to comparing samples to centroids, the Euclidean distance between the centroids C_i and C_j was computed using the following formula:

$$d(C_i, C_j) = \sqrt{(C_{i1} - C_{j1})^2 + (C_{i2} - C_{j2})^2 + (C_{i3} - C_{j3})^2}, \quad (6)$$

This step is particularly important as it provides a measure of how similar or different the classes are from each other in the feature space. Large distances between centroids suggest that the corresponding classes are well-separated, while smaller distances may indicate some overlap or similarity between classes. To facilitate the interpretation of these distances, bar plots were generated for each class as shown in

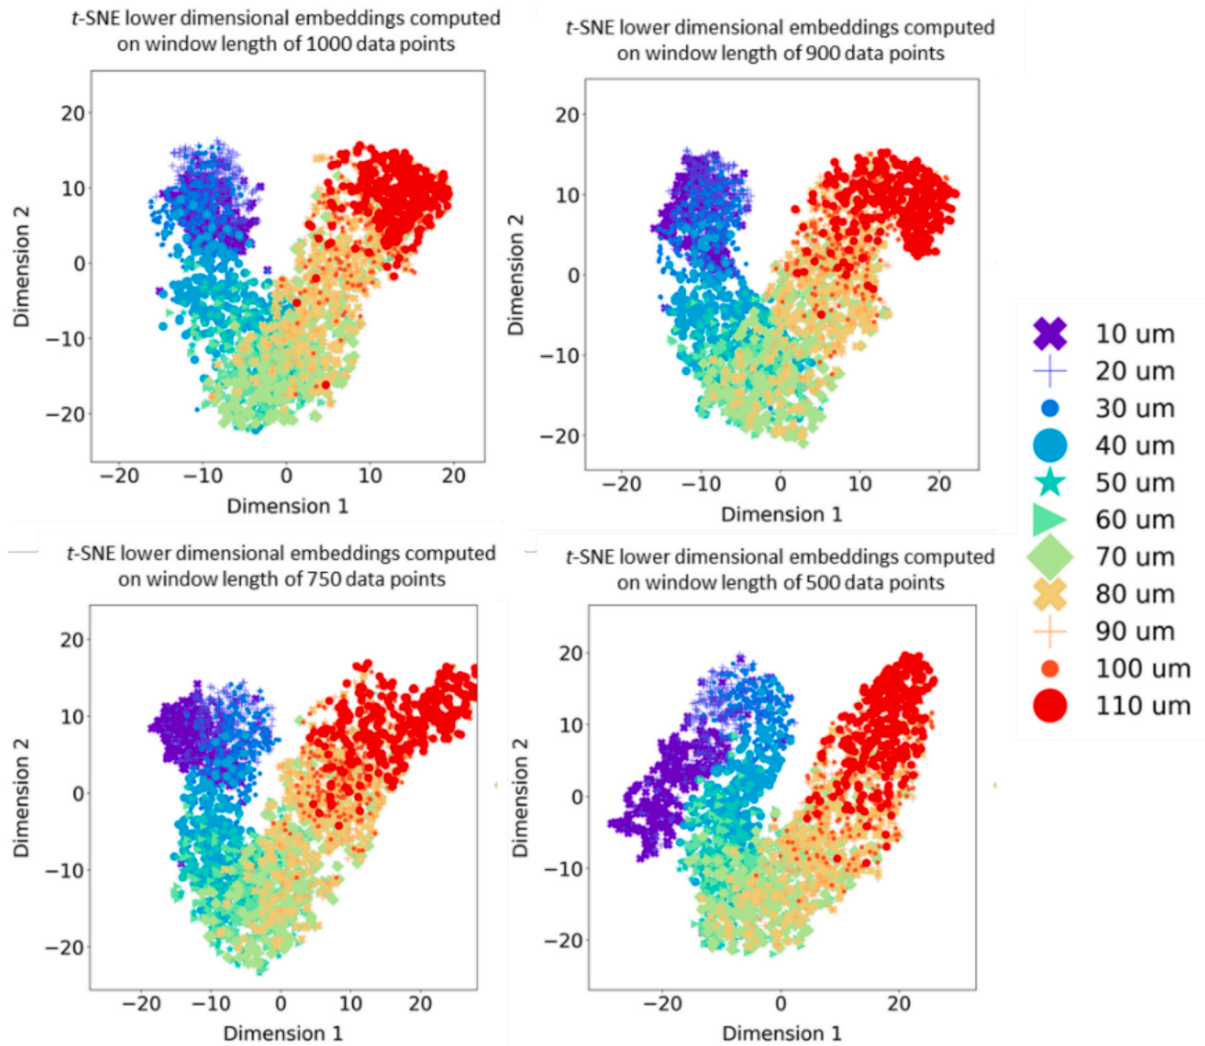


Fig. 15. t-SNE lower embedding on four different time-scales showing similarity.

Fig. 16 on datasets with a window length of 1000 data points. In each plot, the bars represent the distances between the centroid of the class under consideration and the centroids of the other classes. The distance to the class's own centroid was, by definition, zero. The remaining bars were color-coded to reflect the distances to the other classes. Additionally, error bars were added to the plots to represent the variability of the distances between the class centroid and individual samples within the class. These error bars were calculated as the standard deviation of the distances between the centroid and the samples of the class, providing a measure of how tightly or loosely the samples are clustered around their centroid.

The visualization allowed us to identify which classes were more similar or dissimilar based on their distances. Classes with large centroid distances were clearly separated in the feature space, while those with smaller distances exhibited more overlap. The neighbouring classes had much similarity therefore their distances appeared smaller. This method of centroid-based analysis provided an intuitive and effective means of evaluating class separability and structure, helping to understand the behavior of the dataset and the relationships between classes. Such analysis is crucial for tasks like classification, where well-separated classes can lead to higher predictive accuracy, and for clustering, where understanding the distribution of data points relative to class centroids can reveal the underlying structure of the data. The evolution of these centroid distances not only reinforces the reliability of the distance metric but also offers a more complex understanding of how

different signal representations behave relative to each other over time, further validating this approach as a powerful tool for anomaly detection and classification in layer thickness differentiation. This classification process leverages the learned representations and the spatial relationships encoded within the t-SNE embedding to make informed decisions about the depth of the powder in the material.

To quantitatively evaluate the discriminative quality of the learned embeddings of the TCN encoder (f_{θ}), the task was formulated as anomaly detection by designating the 30 μm thickness class as the nominal condition and treating all other thicknesses as anomalies. This binary classification setup assessed the effectiveness of the learned representations in distinguishing normal from anomalous layer thickness patterns. A centroid was computed for the 30 μm class in the latent space using 60 % of the training dataset and 10 % of the validation dataset, with ground truths considered for evaluation. Classification was performed on the remaining 30 % of the test set, which was used neither for TCN training or centroid calculation, based on Euclidean distance from the 30 μm centroid. The 30 μm case was chosen only as a representative example, and the same formulation can be applied to other nominal thicknesses. To establish a robust separation criterion, the threshold was defined as the 95th percentile of centroid distance values from the nominal training data. Since the proposed model is fundamentally designed to operate across variable time scales—a key novelty of the framework—this threshold was first computed for four signal window lengths (1000, 900, 750, and 500 samples), and their average yielded a final global threshold

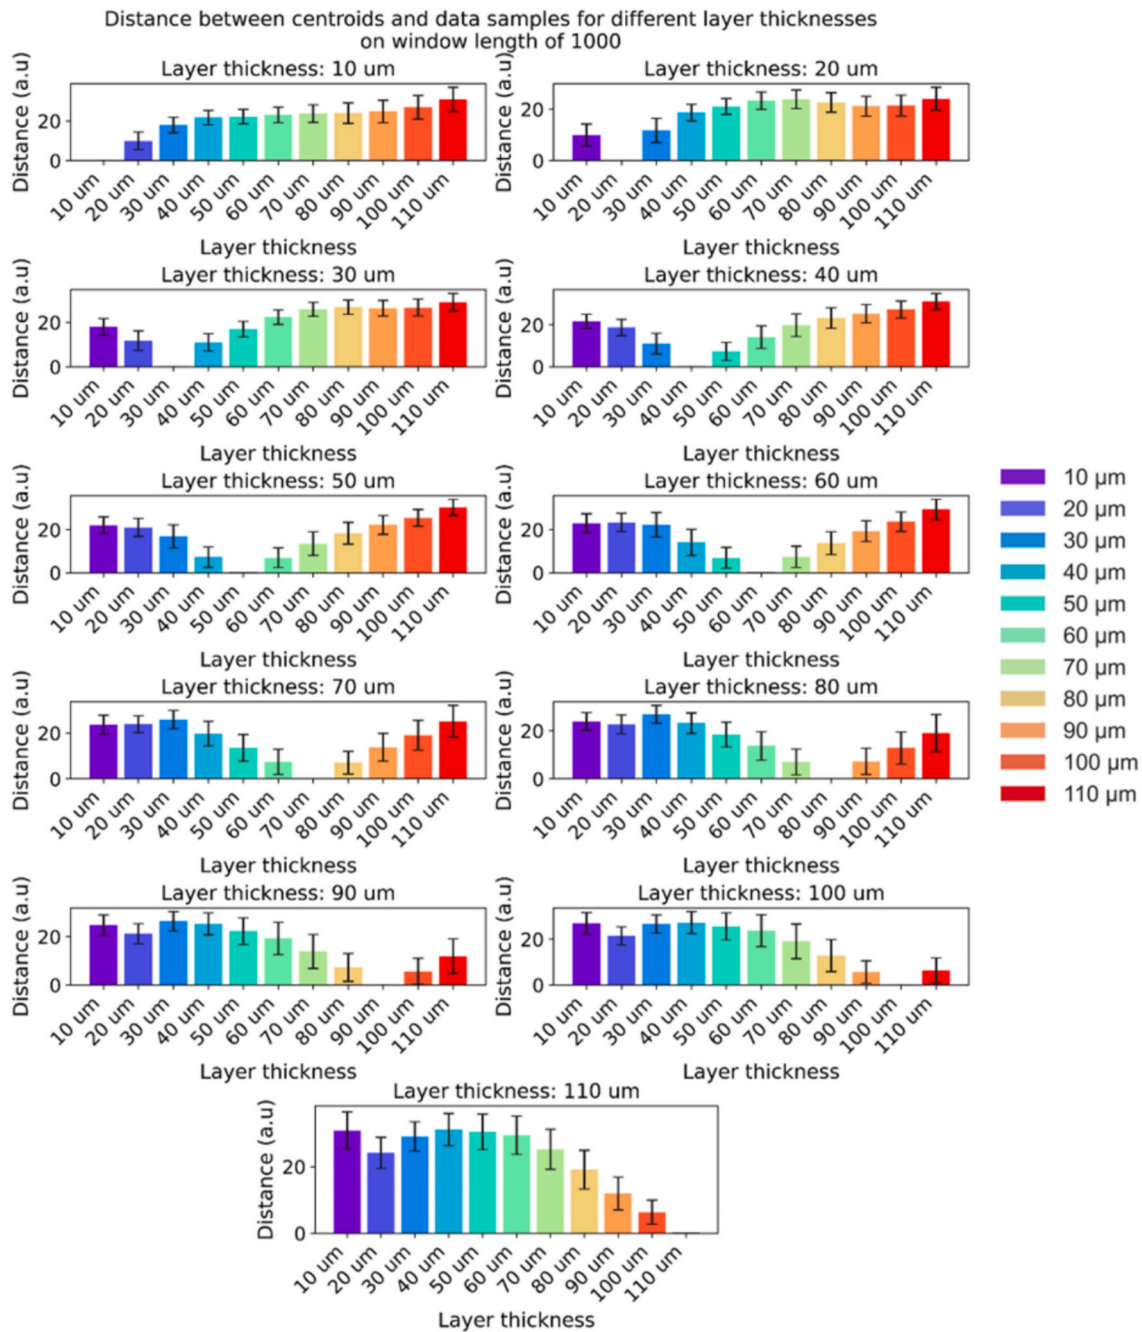


Fig. 16. Distance between centroids and data samples for different layer thicknesses with a window length of 1000.

of 7.8. As shown in Fig. 17, the distance distributions across these window lengths consistently exhibit strong separability, resulting in reliable classification accuracy.

In this setup, the model achieved ~ 90.5 % classification accuracy, offering a clear quantitative metric that complements the *t*-SNE and centroid-based analyses. Most classification errors involved neighboring classes (e.g., 20 μm and 40 μm) being misidentified as normal—an outcome that aligns with the physical difficulty of tightly controlling powder layer thickness in LPBF. It is important to note that the layer thickness distribution forms a continuous space, ranging from 10 μm to 110 μm, which inherently complicates binary classification boundaries in anomaly detection tasks. These results confirm that the self-supervised TCN model effectively captures latent thermal patterns linked to layer thickness variation, enabling robust anomaly detection and binary classification despite the absence of explicit label supervision

during training.

5. Discussion

The self-supervised TCN model developed in this work offers significant potential for integration into closed-loop LPBF control systems. By continuously inferring patterns from pyrometer data sampled at 100 kHz, the model can make decisions within a few milliseconds, depending on the window length. A key novelty of the proposed framework is its robustness to input signal length, allowing it to operate effectively on varying temporal data length without performance degradation, crucial for real-time responsiveness in dynamic scan conditions. If a deviation from the nominal layer thickness is detected, autonomous corrective actions such as adjusting laser power, modifying scan speed, or halting the build can be initiated. Our model is intentionally designed to learn




Fig. 17. Distance-based anomaly detection using latent embeddings from the self-supervised TCN model across four window lengths on the test dataset.

from the thermal emission signatures that arise due to variations in layer thickness, rather than attempting to isolate abstract geometric thickness values. At very thin layers (e.g., 10 μm), the melt pool absorbs more energy per unit mass, often leading to localized over-melting or keyholing. Conversely, at thicker layers (e.g., 110 μm), the same energy input becomes insufficient, resulting in lack-of-fusion or under-melting. Rather than introducing confounding effects, these phenomena represent a realistic and physically coupled relationship between layer thickness deviations and melt pool dynamics in LPBF. This intrinsic coupling between geometric variation and the material's thermophysical response enables the model to associate learned representations with meaningful process indicators—such as keyhole porosity and lack of fusion—based solely on their thermal emission patterns.

As previously discussed in section 4, the overlap in the latent space representation of the TCN for closely spaced layer thicknesses based on optical signatures highlights a critical limitation in sensor resolution. This overlap suggests that the current sensor setup can only effectively distinguish layer thickness variations greater than twenty microns. As layer thickness is a key variable influencing the quality of the LPBF process, the inability to detect finer variations could hinder the accurate monitoring and control of the process. Achieving finer differentiation between layers would require higher-resolution sensors that can capture smaller increments in layer thickness. This limitation brings to light the broader challenge of sensor accuracy in LPBF, where even minute variations in process parameters can lead to significant differences in build quality. For deviations less than 20 μm , higher-resolution sensing modalities—such as visible-spectrum cameras or coaxial high-resolution imaging systems—could provide enhanced spatial granularity. These imaging techniques are particularly well suited for static or dormant conditions, such as capturing a snapshot of the powder bed prior or after laser exposure. However, for dynamic deviations that occur in real time—such as denudation during melting, keyholing in thin layers, or

lack of fusion in thicker layers, the high temporal resolution of pyrometry (100 kHz) offers a key advantage. It allows for continuous monitoring of fast transient melt pool dynamics that may not be detectable with lower-frame-rate imaging systems. Sensor fusion—combining pyrometry with high-resolution imaging—is a promising direction for improving both spatial precision and monitoring robustness in LPBF systems. Also, there is a growing need to explore sensor fusion approaches for continuous monitoring. Integrating additional sensors—such as those that monitor acoustic emissions or detect photonic signals at different wavelengths—could enhance the resolution and provide complementary insights. These data sources may capture different facets of the process zone, offering a more holistic view of the complex interactions during the LPBF process. Fusing data from multiple sources makes it possible to differentiate layer thicknesses with greater precision and detect subtle anomalies or variations that might otherwise go unnoticed. Moreover, sensor fusion could be particularly beneficial in identifying transitions or critical zones within the build process, where changes in layer thickness or material properties may affect part quality. By leveraging a combination of optical, acoustic, and other sensor modalities, the monitoring system could detect deviations early and prevent defects, such as porosity or incomplete fusion, which are often challenging to identify using a single sensor modality.

The TCN backbone model developed for this study was trained using datasets in which layer thickness was the sole variable. This controlled environment allowed a focused investigation of the signal dynamics associated with different layer thicknesses. The TCN learned these dynamics effectively, as the latent space representation shows meaningful differentiation for layer thickness variations above twenty microns. However, in real-world LPBF applications, process signals can be influenced by various factors beyond layer thickness, such as alloy composition, laser power, and scan speed. Variations in alloy composition, in particular, can introduce complexity in the signal profiles, as

different materials exhibit unique optical and acoustic emission characteristics during laser interaction. This variability poses a challenge and an opportunity to extend the model's applicability. Based on the TCN backbone, the proposed framework is inherently adaptable and can be trained on other forms of 1D process data beyond optical signatures. For example, acoustic emissions, which capture the laser regimes generated during the process, offer valuable insights into material behaviour and defect formation. Similarly, photodetector outputs at different wavelengths can provide information on temperature gradients and melt pool dynamics, while laser back-reflection intensity offers clues about surface roughness and reflectivity changes. By integrating these diverse data streams, the framework can be made more robust and capable of monitoring various aspects of the LPBF process in real time. Furthermore, the proposed model is not limited to a specific machine or sensor configuration. LPBF machines often differ slightly regarding sensor placement, calibration, and environmental conditions. These variations can introduce inconsistencies in the data collected from identical processes on different machines. The adaptability of the TCN model allows it to handle such variations, making it an effective solution for monitoring LPBF processes across different machines or setups without requiring extensive retraining.

To assess the feasibility of real-time deployment, the end-to-end latency of the proposed monitoring pipeline was quantified. Pyrometer signals are sampled at 100 kHz, and each inference window spans 1000 samples, corresponding to a 10 ms temporal segment. On an NVIDIA RTX Titan GPU, the combined latency for preprocessing (including normalization and optional augmentation), inference through the TCN encoder, and post-inference classification remains under 1.5 ms. The post-inference classification step based on centroid proximity adds negligible overhead (~ 0.2 ms). As illustrated in Fig. 18, the per-window inference time averages between 0.5–0.6 ms across all dataset splits, with worst-case bounds below 1.2 ms. The intervals shown in the histograms correspond to the standard deviation across multiple runs. These values fall well within typical LPBF scan durations, which range from tens to hundreds of milliseconds per hatch line. This confirms the method's suitability for integration into real-time control loops at the vector or sub-layer level, where responsiveness and low-latency signal processing are critical.

While the model shows promising results in differentiating layer thicknesses under controlled conditions, its generalizability across complex geometries and scan vectors has yet to be fully validated. Intricate geometries, especially those with fine features or high aspect ratios, introduce additional challenges for process monitoring. These geometries often result in irregular scan patterns, varying heat dissipation rates, and non-uniform material melting, all of which can affect the

process emissions. Therefore, testing the model's ability to generalize across these challenging cases is essential to ensure its robustness and reliability in real-world applications. Another key aspect that must be addressed is the model's ability to learn from anomaly regimes. In practical additive manufacturing settings, sub-optimal process parameters—such as incorrect laser power, excessive scan speed, or improper layer thickness—are generally avoided. However, exposing the model to data from these anomaly regimes during training is crucial. This exposure allows the model to learn the boundaries of optimal and suboptimal process conditions, vital for identifying transition zones where defects are likely to occur. Transition zones, such as those between different layer thicknesses or material compositions, are often where issues like porosity, cracking, or delamination originate. By training the model on data from these regimes, we can enhance its ability to detect anomalies early, even in situations that deviate from the norm. Understanding these transition zones through model training offers profound insights into process physics, potentially guiding future process optimizations. This knowledge can inform decisions about adjusting process parameters dynamically to prevent defects before they manifest, ultimately improving the quality and consistency of LPBF-produced parts. As a result, this approach ensures that the model remains robust across a wide range of operating conditions and helps develop a more intelligent, adaptive LPBF process control system capable of maintaining high-quality outputs under varying scenarios.

6. Conclusion

This study introduces a co-axial photodetector-based self-supervised learning approach for real-time layer thickness identification in the LPBF process. By leveraging the high temporal resolution of pyrometers to capture optical emissions from the process zone, we demonstrated the feasibility of accurately monitoring and identifying layer thickness variations. The proposed TCN model, enhanced with dynamic data randomization techniques, successfully learned to differentiate layer thicknesses and their associated anomalies. Our self-supervised methodology addresses several key challenges in LPBF process monitoring by eliminating the need for labelled data and overcoming the constraints of fixed-length input signals.

Analysis of the model's performance through *t*-SNE visualizations revealed distinct clusters corresponding to different layer thicknesses, confirming its ability to capture nuanced differences in optical signals. The study also highlights that the current sensor resolution limits reliable detection to thickness variations above ~ 20 μm , underscoring the need for higher-resolution sensors or sensor fusion strategies to capture finer distinctions. This highlights the need for higher-resolution sensors or sensor fusion techniques to improve precision in future implementations. The centroid-based distance analysis further demonstrated the stability and robustness of the model across different window lengths. Despite slight increases in variance for shorter windows, centroids remained well-separated, validating the model's capacity for layer thickness classification. Additionally, the TCN model's adaptability across multiple 1D data sources suggests its potential for integration with other sensor modalities, such as acoustic emissions or laser back-reflection data. This extension would enhance real-time monitoring capabilities and ensure greater flexibility to accommodate variations in sensor placement across different LPBF machines, thus broadening its industrial applicability.

Looking ahead, future work will focus on several key areas. First, testing and validating the model's general ability on more complex geometries and scan vectors is crucial to ensure its robustness across various LPBF configurations. Second, incorporating data from anomaly regimes, such as suboptimal laser power or excessive scan speed, will enhance the model's ability to detect process anomalies and transition zones, improving defect detection and process optimization strategies. Furthermore, exploring sensor fusion by integrating acoustic and photonic sensors will improve layer thickness resolution and provide a more

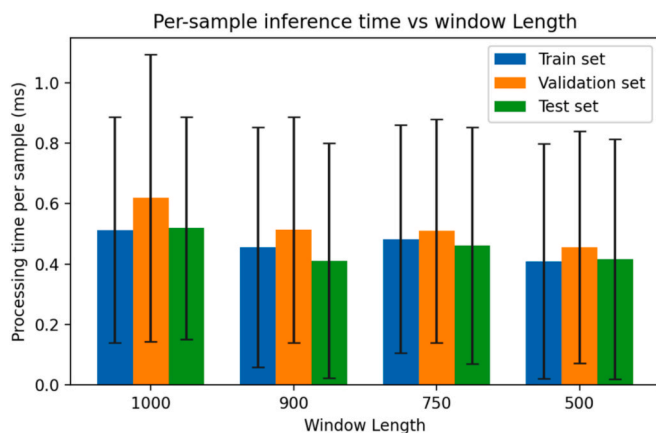


Fig. 18. Inference latency of the TCN network for varying signal window lengths, with bars showing means and intervals denoting one standard deviation.

holistic understanding of the LPBF process. These advancements will contribute to developing an intelligent, adaptive LPBF process control system capable of consistently delivering high-quality outputs under diverse conditions. For those interested, the data and codes for this study can be accessed in the following repositories (<https://doi.org/10.5281/zenodo.11101714>, <https://gitlab.utu.fi/vpsora/Additive-Manufacturing-Pyrometer-Variable-Timescales.git>).

CRedit authorship contribution statement

Bariş Kavas: Writing – original draft, Methodology, Formal analysis. **Roland Axel Richter:** Writing – review & editing. **Michael R. Tucker:** Writing – review & editing. **Vigneshwara Pandiyan:** Writing – review & editing, Writing – original draft, Visualization, Supervision, Methodology, Investigation, Data curation, Conceptualization.

Declaration of competing interest

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

Data availability

As strong advocates of open science and the democratization of AI, we have made both the data and code used in this study publicly available. The repository link is provided in the manuscript

References

- [1] S. Chowdhury, N. Yadaiah, C. Prakash, S. Ramakrishna, S. Dixit, L.R. Gupta, D. Buddhi, Laser powder bed fusion: a state-of-the-art review of the technology, materials, properties & defects, and numerical modelling, *J. Mater. Res. Technol.* 20 (2022) 2109–2172.
- [2] J. Elambasseril, J. Rogers, C. Wallbrink, D. Munk, M. Leary, M. Qian, Laser powder bed fusion additive manufacturing (LPBF-AM): the influence of design features and LPBF variables on surface topography and effect on fatigue properties, *Crit. Rev. Solid State Mater. Sci.* 48 (1) (2023) 132–168.
- [3] J.V. Gordon, S.P. Narra, R.W. Cunningham, H. Liu, H. Chen, R.M. Suter, J.L. Beuth, A.D. Rollett, Defect structure process maps for laser powder bed fusion additive manufacturing, *Addit. Manuf.* 36 (2020) 101552.
- [4] L. Thijs, F. Verhaeghe, T. Craeghs, J. Van Humbeeck, J.-P. Kruth, A study of the microstructural evolution during selective laser melting of Ti–6Al–4V, *Acta Mater.* 58 (9) (2010) 3303–3312.
- [5] X. Tan, Y. Kok, Y.J. Tan, M. Descroins, D. Mangelinck, S.B. Tor, K.F. Leong, C. K. Chua, Graded microstructure and mechanical properties of additive manufactured Ti–6Al–4V via electron beam melting, *Acta Mater.* 97 (2015) 1–16.
- [6] N. Nudelis, P. Mayr, Defect-based analysis of the laser powder bed fusion process using X-ray data, *Int. J. Adv. Manuf. Technol.* 123 (9) (2022) 3223–3232.
- [7] N.D. Dejene, H.G. Lemu, Current status and challenges of powder bed fusion-based metal additive manufacturing: literature review, *Metals* 13 (2) (2023) 424.
- [8] S. Jayasinghe, P. Paoletti, C. Sutcliffe, J. Dardis, N. Jones, P.L. Green, Automatic quality assessments of laser powder bed fusion builds from photodiode sensor measurements, *Prog. Addit. Manuf.* (2021) 1–18.
- [9] I.S. Ramírez, F.P.G. Márquez, M. Papaalias, Review on additive manufacturing and non-destructive testing, *J. Manuf. Syst.* 66 (2023) 260–286.
- [10] J. Rao, S. Leong Sing, P. Liu, J. Wang, H. Sohn, Non-destructive testing of metal-based additively manufactured parts and processes: a review, *Virtual and Physical Prototyping* 18 (1) (2023) e2266658.
- [11] V. Pandiyan, A. Baganis, R. Axel Richter, R. Wröbel, C. Leinenbach, Qualify-as-you-go: sensor fusion of optical and acoustic signatures with contrastive deep learning for multi-material composition monitoring in laser powder bed fusion process, *Virtual and Physical Prototyping* 19 (1) (2024) e2356080.
- [12] V. Pandiyan, G. Masinelli, N. Claire, T. Le-Quang, M. Hamidi-Nasab, C. de Formanoir, R. Esmaeilzadeh, S. Goel, F. Marone, R. Logé, Deep learning-based monitoring of laser powder bed fusion process on variable time-scales using heterogeneous sensing and operando X-ray radiography guidance, *Addit. Manuf.* 58 (2022) 103007.
- [13] Y. Cai, J. Xiong, H. Chen, G. Zhang, A review of in-situ monitoring and process control system in metal-based laser additive manufacturing, *J. Manuf. Syst.* 70 (2023) 309–326.
- [14] T. Özel, A review on in-situ process sensing and monitoring systems for fusion-based additive manufacturing, *International Journal of Mechatronics and Manufacturing Systems* 16 (2–3) (2023) 115–154.
- [15] K. Gutknecht, M. Cloots, R. Sommerhuber, K. Wegener, Mutual comparison of acoustic, pyrometric and thermographic laser powder bed fusion monitoring, *Mater. Des.* 210 (2021) 110036.
- [16] B. Lane, E. Whitenon, S. Moylan, Multiple sensor detection of process phenomena in laser powder bed fusion, *Thermosense: thermal Infrared applications XXXVIII, SPIE* (2016) 20–28.
- [17] M. Doubenskaia, M. Pavlov, S. Grigoriev, E. Tikhonova, I. Smurov, Comprehensive optical monitoring of selective laser melting, *Journal of Laser Micro Nanoengineering* 7 (3) (2012) 236.
- [18] X. Peng, L. Kong, H. An, G. Dong, A review of in situ defect detection and monitoring technologies in selective laser melting, 3D printing and additive manufacturing 10(3) (2023) 438–466.
- [19] J.A. Mitchell, T.A. Ivanoff, D. Dagle, J.D. Madison, B. Jared, Linking pyrometry to porosity in additively manufactured metals, *Addit. Manuf.* 31 (2020) 100946.
- [20] J.-B. Forien, N.P. Calta, P.J. DePond, G.M. Guss, T.T. Roehling, M.J. Matthews, Detecting keyhole pore defects and monitoring process signatures during laser powder bed fusion: a correlation between in situ pyrometry and ex situ X-ray radiography, *Addit. Manuf.* 35 (2020) 101336.
- [21] Y. Gawade, G. Galkin, Y. Guo, W.G. Guo, Quantifying and modeling overheating using 3D pyrometry map in powder bed fusion, *Manuf. Lett.* 33 (2022) 880–892.
- [22] J.-B. Forien, P.J. DePond, G.M. Guss, B.H. Jared, J.D. Madison, M.J. Matthews, Effect of laser power on roughness and porosity in laser powder bed fusion of stainless steel 316L alloys measured by X-ray tomography: Paper presented at the Symposium “Tomographic and Radiographic Imaging with Synchrotron X-rays and Neutrons” of the MSE 2018, 26–28 September 2018, Darmstadt, Germany, *International Journal of Materials Research* 111(1) (2020) 47–54.
- [23] R. Yavari, A. Riensche, E. Tekerek, L. Jacquemetton, H. Halliday, M. Vandever, A. Tenequer, V. Perumal, A. Kontsos, Z. Smoqi, Digitally twinned additive manufacturing: Detecting flaws in laser powder bed fusion by combining thermal simulations with in-situ melt pool sensor data, *Mater. Des.* 211 (2021) 110167.
- [24] M. Ma, Z. Wang, M. Gao, X. Zeng, Layer thickness dependence of performance in high-power selective laser melting of 1Cr18Ni9Ti stainless steel, *J. Mater. Process. Technol.* 215 (2015) 142–150.
- [25] Y. Mahmoodkhani, U. Ali, S. Imani Shahabad, A. Rani Kasinathan, R. Esmaeilzadeh, A. Keshavarzkermani, E. Marzbanrad, E. Toyserkani, On the measurement of effective powder layer thickness in laser powder-bed fusion additive manufacturing of metals, *Prog. Addit. Manuf.* 4 (2) (2019) 109–116.
- [26] L. Xiao, M. Lu, H. Huang, Detection of powder bed defects in selective laser sintering using convolutional neural network, *Int. J. Adv. Manuf. Technol.* 107 (5) (2020) 2485–2496.
- [27] A. Hehr, M. Norfolk, D. Kominsky, A. Boulanger, M. Davis, P. Boulware, Smart Build-Plate for Metal Additive Manufacturing Processes, *Sensors* 20 (2) (2020) 360.
- [28] M.A. Ansari, A. Crampton, S. Parkinson, A Layer-Wise Surface deformation defect detection by convolutional neural networks in laser powder-bed fusion images, *Materials* 15 (20) (2022) 7166.
- [29] S.C. Subedi, A. Shahba, M. Thevamaran, D.J. Thoma, K. Suresh, Towards the optimal design of support structures for laser powder bed fusion-based metal additive manufacturing via thermal equivalent static loads, *Addit. Manuf.* 57 (2022) 102956.
- [30] S. Brenner, M. Moser, L. Strauß, V. Nedeljkovic-Groha, G. Löwisch, Recoater crashes during powder bed fusion of metal with laser beam: simulative prediction of interference and experimental evaluation of resulting part quality, *Prog. Addit. Manuf.* 8 (4) (2023) 759–768.
- [31] X. Lu, M. Chiumentti, M. Cervera, M. Slimani, I. Gonzalez, Recoater-induced distortions and build failures in selective laser melting of thin-walled Ti6Al4V Parts, *Journal of Manufacturing and Materials Processing* 7 (2) (2023) 64.
- [32] J.-P. Kruth, J. Deckers, E. Yasa, R. Wauthlé, Assessing and comparing influencing factors of residual stresses in selective laser melting using a novel analysis method, *Proceedings of the Institution of Mechanical Engineers, Part B* 226(6) (2012) 980–991.
- [33] M.F. Zaeh, G. Branner, Investigations on residual stresses and deformations in selective laser melting, *Prod. Eng.* 4 (1) (2010) 35–45.
- [34] T. Mukherjee, T. DebRoy, Mitigation of lack of fusion defects in powder bed fusion additive manufacturing, *J. Manuf. Process.* 36 (2018) 442–449.
- [35] S.K. Nayak, S.K. Mishra, C.P. Paul, K.S. Bindra, Effect of higher layer thickness on laser powder bed fusion built single tracks of Ni-Cr-Fe-Nb-Mo alloy, *Journal of Micromanufacturing* 4 (2) (2021) 169–178.
- [36] R.J. Williams, C.M. Davies, P.A. Hooper, In situ monitoring of the layer height in laser powder bed fusion, *Mater. Des. Process. Commun.* 3 (6) (2021) e173.
- [37] T. Liu, C.S. Lough, H. Sehhat, Y.M. Ren, P.D. Christofides, E.C. Kinzel, M.C. Leu, In-situ infrared thermographic inspection for local powder layer thickness measurement in laser powder bed fusion, *Addit. Manuf.* 55 (2022) 102873.
- [38] T. Liu, E.C. Kinzel, M.C. Leu, In-situ lock-in thermographic measurement of powder layer thermal diffusivity and thickness in laser powder bed fusion, *Addit. Manuf.* 74 (2023) 103726.
- [39] B. Kavas, E.C. Balta, M. Tucker, A. Rupenyan, J. Lygeros, M. Bambach, Layer-to-layer closed-loop feedback control application for inter-layer temperature stabilization in laser powder bed fusion, *Addit. Manuf.* 78 (2023) 103847.
- [40] Y. Wang, W. Guo, H. Li, Y. Xie, J. Shi, Z. Liang, P. Han, S. Li, H. Zhang, Nano-scale microstructural evolution and mechanical property enhancement mechanism during crack inhibition in nickel-based superalloys fabricated by laser powder bed fusion, *Addit. Manuf.* 100 (2025) 104685.
- [41] Y. Zhang, M.H.R. Dehkordi, M.J. Kholoud, H. Azimy, Z. Li, M. Akbari, Numerical modeling of the temperature distribution and melt flow in dissimilar fiber laser welding of duplex stainless steel 2205 and low alloy steel, *Opt. Laser Technol.* 174 (2024) 110575.
- [42] S.K. Everton, M. Hirsch, P. Stravroulakis, R.K. Leach, A.T. Clare, Review of in-situ process monitoring and in-situ metrology for metal additive manufacturing, *Mater. Des.* 95 (2016) 431–445.

- [43] G. Tapia, A. Elwany, A review on process monitoring and control in metal-based additive manufacturing, *J. Manuf. Sci. Eng.* 136 (6) (2014).
- [44] M. Grasso, A. Demir, B. Previtali, B. Colosimo, In situ monitoring of selective laser melting of zinc powder via infrared imaging of the process plume, *Rob. Comput. Integr. Manuf.* 49 (2018) 229–239.
- [45] Z. Yan, W. Liu, Z. Tang, X. Liu, N. Zhang, M. Li, H. Zhang, Review on thermal analysis in laser-based additive manufacturing, *Opt. Laser Technol.* 106 (2018) 427–441.
- [46] V. Pandiyan, R. Drissi-Daoudi, S. Shevchik, G. Masinelli, R. Logé, K. Wasmer, Analysis of time, frequency and time-frequency domain features from acoustic emissions during Laser Powder-Bed fusion process, *Procedia CIRP* 94 (2020) 392–397.
- [47] R. Drissi-Daoudi, V. Pandiyan, R. Logé, S. Shevchik, G. Masinelli, H. Ghasemi-Tabasi, A. Parrilli, K. Wasmer, Differentiation of materials and laser powder bed fusion processing regimes from airborne acoustic emission combined with machine learning, *Virtual and Physical Prototyping* 17 (2) (2022) 181–204.
- [48] R. Esmailzadeh, V. Pandiyan, S. Van Petegem, M. Van der Meer, M.H. Nasab, C. de Formanoir, J. Jhabvala, C. Navarre, L. Schlenger, R. Richter, Acoustic emission signature of martensitic transformation in laser powder bed fusion of Ti6Al4V-Fe, supported by operando X-ray diffraction, *Addit. Manuf.* 96 (2024) 104562.
- [49] M. Khanzadeh, S. Chowdhury, M. Maruffuzaman, M.A. Tschopp, L. Bian, Porosity prediction: Supervised-learning of thermal history for direct laser deposition, *J. Manuf. Syst.* 47 (2018) 69–82.
- [50] M.A. Ansari, A. Crampton, R. Garrard, B. Cai, M. Attallah, A Convolutional Neural Network (CNN) classification to identify the presence of pores in powder bed fusion images, *Int. J. Adv. Manuf. Technol.* 120 (7) (2022) 5133–5150.
- [51] C. Gobert, E.W. Reutzel, J. Petrich, A.R. Nassar, S. Phoha, Application of supervised machine learning for defect detection during metallic powder bed fusion additive manufacturing using high resolution imaging, *Addit. Manuf.* 21 (2018) 517–528.
- [52] M. Khanzadeh, W. Tian, A. Yadollahi, H.R. Doude, M.A. Tschopp, L. Bian, Dual process monitoring of metal-based additive manufacturing using tensor decomposition of thermal image streams, *Addit. Manuf.* 23 (2018) 443–456.
- [53] I. Okaro, S. Jayasinghe, C. Sutcliffe, K. Black, P. Paoletti, P. Green, Automatic Fault Detection for Selective Laser Melting using Semi-Supervised Machine Learning, *Addit. Manuf.* (2018) 1–35.
- [54] G. Goh, S. Sing, W. Yeong, A review on machine learning in 3D printing: applications, potential, and challenges, *Artif. Intell. Rev.* 54 (1) (2021) 63–94.
- [55] W. Xing, X. Chu, T. Lyu, C.-G. Lee, Y. Zou, Y. Rong, Using convolutional neural networks to classify melt pools in a pulsed selective laser melting process, *J. Manuf. Process.* 74 (2022) 486–499.
- [56] O. Kwon, H.G. Kim, M.J. Ham, W. Kim, G.-H. Kim, J.-H. Cho, N.I. Kim, K. Kim, A deep neural network for classification of melt-pool images in metal additive manufacturing, *J. Intell. Manuf.* 31 (2) (2020) 375–386.
- [57] J. Li, Q. Zhou, L. Cao, Y. Wang, J. Hu, A convolutional neural network-based multi-sensor fusion approach for in-situ quality monitoring of selective laser melting, *J. Manuf. Syst.* 64 (2022) 429–442.
- [58] A. Gaikwad, B. Giera, G.M. Guss, J.-B. Forien, M.J. Matthews, P. Rao, Heterogeneous sensing and scientific machine learning for quality assurance in laser powder bed fusion—a single-track study, *Addit. Manuf.* 36 (2020) 101659.
- [59] B. Yuan, B. Giera, G. Guss, S. McMains, Semi-Supervised Learning for Selective Laser Melting Monitoring via In-Situ Videos, Lawrence Livermore National Lab. (LLNL), Livermore, CA (United States), 2018.
- [60] V. Pandiyan, R. Drissi-Daoudi, S. Shevchik, G. Masinelli, T. Le-Quang, R. Logé, K. Wasmer, Semi-supervised monitoring of Laser powder bed fusion process based on acoustic emissions, *Virtual and Physical Prototyping* 16 (4) (2021) 481–497.
- [61] R. McCann, M.A. Obeidi, C. Hughes, É. McCarthy, D.S. Egan, R.K. Vijayaraghavan, A.M. Joshi, V.A. Garzon, D.P. Dowling, P.J. McNally, In-situ sensing, process monitoring and machine control in Laser Powder Bed Fusion: a review, *Addit. Manuf.* 45 (2021) 102058.
- [62] D. Guillen, S. Wahlquist, A. Ali, Critical Review of LPBF Metal Print Defects Detection: Roles of Selective Sensing Technology, *Appl. Sci.* 14 (15) (2024) 6718.
- [63] P. Becker, C. Roth, A. Rönau, R. Dillmann, Porosity Detection in Powder Bed Fusion Additive Manufacturing with Convolutional Neural Networks.
- [64] V. Pandiyan, R. Drissi-Daoudi, S. Shevchik, G. Masinelli, T. Le-Quang, R. Logé, K. Wasmer, Deep transfer learning of additive manufacturing mechanisms across materials in metal-based laser powder bed fusion process, *J. Mater. Process. Technol.* 303 (2022) 117531.
- [65] S.A. Shevchik, G. Masinelli, C. Kenel, C. Leinenbach, K. Wasmer, Deep learning for in situ and real-time quality monitoring in additive manufacturing using acoustic emission, *IEEE Trans. Ind. Inf.* 15 (9) (2019) 5194–5203.
- [66] S.A. Shevchik, C. Kenel, C. Leinenbach, K. Wasmer, Acoustic emission for in situ quality monitoring in additive manufacturing using spectral convolutional neural networks, *Addit. Manuf.* 21 (2018) 598–604.
- [67] V. Pandiyan, D. Cui, R.A. Richter, A. Parrilli, M. Leparoux, Real-time monitoring and quality assurance for laser-based directed energy deposition: integrating co-axial imaging and self-supervised deep learning framework, *J. Intell. Manuf.* 36 (2) (2025) 909–933.
- [68] Z.Y. Chua, I.H. Ahn, S.K. Moon, Process monitoring and inspection systems in metal additive manufacturing: Status and applications, *International Journal of Precision Engineering and Manufacturing-Green Technology* 4 (2) (2017) 235–245.
- [69] M. Van Elsen, Complexity of Selective Laser Melting: a new optimisation approach, (2007).
- [70] T.G. Spears, S.A. Gold, In-process sensing in selective laser melting (SLM) additive manufacturing, *Integrating Materials and Manufacturing Innovation* 5 (1) (2016) 16–40.
- [71] D. Gu, Y. Shen, J. Yang, Y. Wang, Effects of processing parameters on direct laser sintering of multicomponent Cu based metal powder, *Mater. Sci. Technol.* 22 (12) (2006) 1449–1455.
- [72] T. Kurzynowski, E. Chlebus, B. Kuźnicka, J. Reiner, Parameters in selective laser melting for processing metallic powders, *High Power Laser Materials Processing: Lasers, Beam Delivery, Diagnostics, and Applications*, International Society for Optics and Photonics (2012) 823914.
- [73] V. Pandiyan, R. Wróbel, R.A. Richter, M. Leparoux, C. Leinenbach, S. Shevchik, Self-Supervised Bayesian representation learning of acoustic emissions from laser powder bed Fusion process for in-situ monitoring, *Mater. Des.* 235 (2023) 112458.