

Accepted Manuscript

Title: Automatic Fault Detection for Laser Powder-Bed Fusion using Semi-Supervised Machine Learning

Author: Ikenna A. Okaro Sarini Jayasinghe Chris Sutcliffe
Kate Black Paolo Paoletti Peter L. Green



PII: S2214-8604(18)30609-2
DOI: <https://doi.org/doi:10.1016/j.addma.2019.01.006>
Reference: ADDMA 626

To appear in:

Received date: 14 August 2018
Revised date: 6 December 2018
Accepted date: 23 January 2019

Please cite this article as: Ikenna A. Okaro, Sarini Jayasinghe, Chris Sutcliffe, Kate Black, Paolo Paoletti, Peter L. Green, Automatic Fault Detection for Laser Powder-Bed Fusion using Semi-Supervised Machine Learning, *Additive Manufacturing* (2019), <https://doi.org/10.1016/j.addma.2019.01.006>

This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

Automatic Fault Detection for Laser Powder-Bed Fusion using Semi-Supervised Machine Learning

Ikenna A. Okaro^{a,b}, Sarini Jayasinghe^b, Chris Sutcliffe^{b,c}, Kate Black^b,
Paolo Paoletti^{a,b}, Peter L. Green^{a,b}

^a*Institute for Risk and Uncertainty, University of Liverpool, Chadwick Building, Peach Street, Liverpool L69 7ZF, United Kingdom*

^b*School of Engineering, University of Liverpool, Harrison Hughes Building, Liverpool, L69 3GH, United Kingdom*

^c*Renishaw Plc, Brooms Road, Stone Business Park, Stone ST15 0SH, United Kingdom*

Abstract

Risk-averse areas such as the medical, aerospace and energy sectors have been somewhat slow towards accepting and applying Additive Manufacturing (AM) in many of their value chains. This is partly because there are still significant uncertainties concerning the quality of AM builds.

This paper introduces a machine learning algorithm for the automatic detection of faults in AM products. The approach is *semi-supervised* in that, during training, it is able to use data from both builds where the resulting components were certified and builds where the quality of the resulting components is unknown. This makes the approach cost efficient, particularly in scenarios where part certification is costly and time consuming.

The study specifically analyses Laser Powder-Bed Fusion (L-PBF) builds. Key features are extracted from large sets of photodiode data, obtained during the building of 49 tensile test bars. Ultimate tensile strength (UTS) tests were then used to categorise each bar as ‘faulty’ or ‘acceptable’. Using a variety of approaches (Receiver Operating Characteristic (ROC) curves and 2-fold cross-validation), it is shown that, despite utilising a fraction of the available certification data, the semi-supervised approach can achieve results comparable to a benchmark case where all data points are labelled. The results show that semi-supervised learning is a promising approach for the

*Corresponding author

Email address: plgreen@liverpool.ac.uk (Peter L. Green)

automatic certification of AM builds that can be implemented at a fraction of the cost currently required.

Keywords: Laser Powder-Bed Fusion, Process Control, Semi-supervised Machine Learning, Randomised Singular Value Decomposition.

1. Introduction

There is a growing demand for efficient manufacturing technologies [1]. Additive Manufacturing (AM) has huge potential in healthcare for custom implants and in aerospace for lightweight designs [2]. However, uncertainties surrounding part quality prevent the full adoption of AM technology in such sectors. Moreover, certification of AM parts is challenging (as faults may occur internal to the parts) and often requires expensive CT scans.

This paper specifically considers Laser Powder-Bed Fusion (L-PBF). L-PBF is an AM technology that has become very popular in recent times due to its ability to produce complex metal geometries, relative to traditional methods. The L-PBF process involves layer-by-layer construction of a build by repeatedly channelling laser beams onto a thin layer of metal powder deposited on a fusion bed [3]. Powder deposition and sintering are repeated until a desired product is made to specification.

For the work detailed herein, data related to the L-PBF process was gathered using high precision photodiodes, which were installed axial to the laser. These sensors were placed behind filters that were designed to eliminate the reflected laser light, thus allowing the reflected light intensity to be measured during the builds. The photodiodes provide process measurements from which, potentially, it may be possible to determine the quality of AM products¹. Understanding the correlations between this data and build quality is a challenging area. However, advances in machine learning have made it possible to create and apply intelligent algorithms to large datasets for decision making [4]. Such algorithms can identify patterns in large data, after

¹Here, the definition of ‘quality’ depends on the product’s intended purpose - it could be used to refer to a build’s surface finish, toughness, ductility etc. In the current paper, ‘quality’ refers to mechanical properties, specifically the ultimate tensile strength of L-PBF builds of tensile test bars.

being trained. The current work is based on the hypothesis that, using large amounts of process measurements from L-PBF machines, machine learning can be used to quickly and cheaply classify the success of L-PBF builds.

Classification algorithms can be broadly categorised as supervised, semi-supervised or unsupervised (for a theoretical review of these methods, [5][6][7] are recommended). With a supervised approach, the algorithm is presented with labelled data - a set of input vectors, each of which is associated with an observed output value (or 'label'). Unsupervised learning can be thought of as finding patterns in only unlabelled data (clustering, for example, is one form of unsupervised learning). With a semi-supervised approach, the user provides some labelled data and some unlabelled data at the same time. The model may then attempt to establish a decision boundary and classifies the data into clusters; based on the characteristics of the provided labelled *and* unlabelled information [8][9].

In the current context, input vectors consist of data that was gathered during AM builds and the labels are used to indicate whether each particular build was 'acceptable' or 'faulty' (in this paper, for example, labels are defined based on ultimate tensile strength values). Consequently, before the application of supervised machine learning, one would have to conduct and certify a large number of AM builds (see [10], for example, where 100s of parts were produced to generate the data needed to train a support vector machine). This procedure would have to be repeated per new type of component or material. However, in many practical applications, completely labelled information is not available [11]. It is more common to find few labelled data and relatively large amounts of unlabelled data. In the current study, for instance, process measurements are generated whenever a component is manufactured, but cost constraints prevent labels from being assigned to most of this data. This study explores how machine learning could help to automatically detect defects in situations where there is a large amount of unlabelled data (builds that were not certified) and a small amount of labelled data (builds that were certified). Furthermore, it illustrates the application of a probabilistic methodology - an important aspect of the approach which allows one to quantify the uncertainties associated with machine-learned assessments of AM builds.

The paper makes 3 main contributions:

1. It is illustrated how a Randomised Singular Value Decomposition can be used to extract key features from large sets of L-PBF process measurements.
2. The feasibility of using machine learning to detect unsuccessful L-PBF builds from process measurements is demonstrated. This highlights how signal-based process monitoring, which is adopted in risk analysis and industrial statistics frameworks [12], could be extended to AM applications.
3. It is shown that, using semi-supervised learning, the number of costly certification experiments associated with such an approach can be significantly reduced.

It is important to note that this study does not aim to draw links between specific L-PBF process parameters and the quality of the resulting builds. Rather, it details a purely data-based approach whereby a machine learning algorithm is used to classify L-PBF build quality based only on the patterns that are contained within sets of photodiode measurements.

The paper is structured as follows; Section 2 discusses current state-of-the-art and highlights the contributions of the paper, Section 3 discusses the semi-supervised model derivation and formulation, Section 4 demonstrates the model using a case study and Section 5 is the conclusion of the work. It is noted that Section 3 is included so that the machine learning approach is not presented as a 'black box'. Section 3 can, however, be skipped by those who are purely interested in the case study.

2. Literature review

This section highlights key relevant contributions before establishing where the current paper fits amongst other literature in the field.

2.1. Key process parameters

Local defects may occur during layer by layer construction of an AM part. The root causes could be traced to improper process parameters, insufficient supports, a non-homogeneous powder deposition, improper heat exchanges

and/or material contaminations [4][13][14]. The effects of process parameters, namely; laser power, scanning speed, hatch spacing, layer thickness and powder temperature, on the tensile strength of AM products are reported in [15][16]. Specifically for L-PBF, it has been shown that four key parameters, namely; part bed temperature, laser power, scan speed, and scan spacing, have significant effect on the mechanical properties and quality of an L-PBF product [17][18][19]. Sensitivity analyses of L-PBF process parameters have revealed that both the scan speed of the laser and scan spacing can be used to facilitate effective improvement in mechanical properties [20].

The type of laser employed determines, to a great extent, the behaviour of the powdered particles during L-PBF processing [21]. This dependence is attributed to the dependence of material laser absorptivity on the wavelength of the laser type used. It has also been discovered that particle size, size distribution, tap density, oxide film thickness, surface chemistry and impurity concentration has little effect on the sintering behaviour of aluminium powders [22]. Debroy et al. [15] pointed out that during laser sintering of metals, the alloying elements vaporise from the surface of the molten pool and, as a result, the surface area-to-volume ratio is one of the crucial factors for determining the magnitude of composition change.

2.2. Monitoring approaches

A significant research effort has been made in recent years to develop monitoring techniques capable of detecting build defects during the build process itself. Such techniques can be broadly classified into two families: *off-axial* and *co-axial*.

Techniques belonging to the first family use sensors placed outside the optical path, and typically use cameras to analyse what is happening on the build plate before and after the layer scans. For example, image segmentation can be used to detect where the outline of the melted region in the layer differ from the nominal position [23]. A combination of cameras and projectors can also be used to reconstruct 3D topography after each layer scan using fringe analysis [24] [25].

Co-axial sensing technologies typically use a combination of pyrometers and cameras, using the optical path of the melting laser to monitor the melt pool. For example, in the setup described in [26], a near-infrared camera

was used to monitor geometry and temperature distribution of the melt pool while a photodiode measured melt pool intensity. Similarly, in [27], a camera and a pyrometers were used to monitor the melt pool, but using an additional co-axial source of illumination. Two pyrometers were used in [28] to capture melt pool emissions at different wavelengths, while a CCD camera measured the image brightness profiles along the transversal and longitudinal directions. An alternative approach, based on interferometry, is described in [29] as well.

For a comprehensive review of different approaches used to monitor the build process and classify defects, interested readers are referred to reference [30].

2.3. Towards feedback control for L-PBF

While much work has been conducted to identify key parameters that affect L-PBF build quality, it can still be difficult to relate this knowledge to the development of effective control strategies. This is particularly evident when one considers developing control strategies for new materials. While proof-of-concept controllers have been developed in [31] [32] (using measurements from high-speed cameras and/or photodiodes to control laser power) and the works [33] [34] detail an approach whereby geometrical accuracy was improved by varying beam offset and scan speed, the adaptability of these methods to new materials can be prohibitively time consuming and/or expensive. In [33], for example, it is stated that ‘*the benchmarking process is time consuming*’ and that ‘*a change of material used will require identification of a new process benchmark as the properties of different materials influence the fabrication parameters in the process*’.

It is worth noting that Finite Element models can, potentially, aid controller development by relating key process parameters to the microstructure of builds (see [35], for example). Unfortunately, these models tend to be very specific in terms of part design and can take a long time to develop and/or implement.

2.4. Machine learning approaches

Through data-based approaches, facilitated by machine learning algorithms, it may be possible to overcome the challenges associated with inferring build quality from knowledge of key parameters and/or the results

of Finite Element models. Work has shown that data-driven methods can, from build data, model how process parameters affect the quality of final parts [36][37][38]. Approaches that utilise build data are advantageous because they provide great opportunities for digitalisation and smart process control, otherwise known as ‘smart manufacturing’ [39][40][41].

Broadly, machine-learning approaches can be categorised as being either ‘supervised’ or ‘unsupervised’. Supervised approaches involve training an algorithm on a set of data, whereby each training point has a ‘label’ attached to it. This label indicates the particular class that the training point belongs to (for example, in the current context, the label could indicate whether a particular set of build data corresponds to a build that had been found to be ‘acceptable’ or ‘faulty’). Supervised algorithms then attempt to infer decision boundaries that separate these classes. Unsupervised approaches, on the other hand, are used to identify key patterns in data that is unlabelled (cluster analysis, for example, is a well-known example of unsupervised learning). The following 2 subsections highlight relevant applications of unsupervised and supervised approaches within the context of L-PBF. Particular attention is given to describing the data acquisition process and/or assumptions involved in each example, as this motivates the use of semi-supervised learning in the current study (Section 2.5).

2.4.1. Unsupervised learning

In [42], to automatically detect local overheating phenomenon, the k-means algorithm was used to cluster features that had been extracted from images (in the visible spectrum) of L-PBF builds. This unsupervised approach clustered the data before further assumptions were used to relate the results to build quality. (It was, for example, assumed that data from ‘normal’ and ‘abnormal’ builds would be best represented by 2 and 3 clusters respectively). The authors of [43] used an unsupervised cluster-based approach to relate melt pool characteristics to build porosity. As with [42], a set of assumptions were then used to relate the clustering results to build quality (specifically, it was assumed that the number of ‘abnormal’ melt pools would be small compared to the number of ‘normal’ melt pools).

In a recent study [44], anomaly detection and classification of SLM specimens was implemented using an unsupervised machine learning algorithm

operating on a training database of image patches. The algorithm functioned well as a post-build analysis tool, allowing a user to identify failure modes and locate regions within a final part that may contain millimeter-scale flaws. However, the algorithm was not designed to classify a mixture of labelled images and unlabelled images simultaneously. The image patches were manually selected from a secondary database; based on a pre-determined rule which clearly distinguished the patches.

2.4.2. *Supervised learning*

Ref. [45] describes a variety of approaches that can be used to infer a relationship between melt pool characteristics and part porosity. To label melt pool data, 3D CT scans were used to empirically locate part defects before algorithm training could begin. In [46] a Gaussian process was used to infer a mapping between laser power and scan speed to part porosity. To facilitate this approach, data was generated by conducting experiments across a grid of laser power and scan speed values before porosity was measured using Archimedes' principle. In [47], a support vector machine was used to classify images of build layers that had been obtained using a high-resolution digital camera. Training data was obtained using 3D CT scans, which were used to identify discontinuities in parts post-build.

It is worth noting that, whilst optical methods have been gaining popularity in recent times, the feasibility of applying supervised learning to acoustic emission data for in-situ quality monitoring has also been investigated. [48], for example, used neural networks to classify features of in-build acoustic signals to one of three quality ratings (defined based on part porosity). To generate data for algorithm training, a SLM build was conducted using different laser scanning parameters which, through a visual analysis, were shown to influence porosity. [49] used acoustic signals to train deep belief networks (a neural network algorithm, sometimes referred to under the banner of 'deep learning'). During data acquisition, process parameters were varied to deliberately induce different types of build flaw, leading to 5 different classes ('balling', 'slight balling', 'normal', 'slight overheating' and 'overheating'). This labelled data was then used to train the parameters of the neural network.

2.5. Contribution of the current work

The works listed above highlight the advantages and disadvantages that can be encountered when implementing both supervised and unsupervised approaches.

Supervised approaches require sufficient quantities of labelled data. Unfortunately, the assignment of labels to data often requires a significant amount of additional resources. [45][47], for example, assigned labels based on the outcomes of CT scans while [46] utilised the results from component porosity tests. To circumvent the requirement for such additional testing, [48] and [49] used pre-existing knowledge regarding the relationship between process parameters and build defects to label data. Such an approach, however, relies on the availability of relatively in-depth knowledge regarding process-defect relationships. This information may be difficult to obtain, particularly when new materials are being analysed.

Unsupervised approaches do not need labelled data and, consequently, are often cheaper to implement. However, the relationship between the results of an unsupervised analysis and build quality has to be built upon an additional set of assumptions. For example, both [42] and [43] had to make assumptions about the number / relative size of the data clusters revealed by their analyses. While the results reported in [42] and [43] are encouraging, it is likely that the validity of these assumptions will come into question if such an approach was used to guarantee build quality for applications in risk-averse disciplines. [44] used an unsupervised approach, but only after image data had been manually selected from a database - a process which, it must be assumed, was fairly time consuming.

In the current work it is suggested that an efficient approach should be able to utilise data that is both labelled *and* unlabelled as, in the author's experience, developing new materials and/or components in L-PBF often leads to a large amount of unlabelled data and a small amount of labelled data. For example, because of the difficulties associated with identifying flaws in AM builds, a set of witness components are usually built alongside the real part. Testing of the witness components then provides a set of labelled data. Clearly, by having more labelled data, it is possible to generate a more reliable analysis of the build-process variability. Testing of the witness components, however, can be expensive and time consuming (particularly

if fatigue tests or CT scans are being utilised). Moreover, the destructive nature of some testing approaches makes it impossible to apply the full set of available testing procedures to each individual witness component. These restrictions lead to many of the witness components being ‘unlabelled’ with regard to certain testing procedures i.e. there exists sets of process measurements for components that were not subject to specific material tests.

The work herein hypothesises that the large amounts of unlabelled data (i.e. process measurements where the final build has not been certified / tested in a particular manner) should not be wasted and, should be analysed alongside the more limited set of labelled data. This *semi-supervised* approach is especially suited to situations where there are few labelled data and much unlabelled data. It therefore has the potential to reduce the number of costly and time consuming certification experiments that are required in the development of machine-learnt models of L-PBF build quality.

3. Model Derivation and Formulation

The proposed semi-supervised method uses a Gaussian Mixture Model (GMM) to classify ‘acceptable’ and ‘faulty’ AM builds. GMMs are often relatively time-efficient as their parameters can be estimated with the Expectation Maximization (EM) algorithm [50],[51] (described briefly in Section 3.2.1). A description of GMMs is covered here for the sake of completeness and to highlight the application of semi-supervised learning in the current context. For more information on GMMs the book [52] is recommended.

Essentially, a GMM algorithm clusters data based on the assumption that each data point is a sample from a mixture of Gaussian distributions, such that the probability distribution over each data point can be described as a weighted sum of Gaussian components [53]. This is elaborated further below, where we first describe how a GMM can be applied to labelled data (supervised learning), before then describing its application to unlabelled data (unsupervised learning). This then helps to establish how a GMM can be used to address situations involving both labelled and unlabelled data.

3.1. Supervised Learning with a GMM

In the following, we use \mathbf{x} to represent an *input vector* - a vector of features that have been extracted from the process measurements of an SLM build. As stated previously, a GMM assumes that each vector, \mathbf{x} , was sampled from a mixture of Gaussian distributions [52] such that

$$p(\mathbf{x}) = \sum_{k=1}^K \pi_k \mathcal{N}(\mathbf{x}; \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k) \quad (1)$$

where $\boldsymbol{\mu} = \{\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_K\}$ represents the means of the Gaussians, $\boldsymbol{\Sigma} = \{\boldsymbol{\Sigma}_1, \dots, \boldsymbol{\Sigma}_K\}$ are covariance matrices of the Gaussians and $\boldsymbol{\pi} = \{\pi_1, \dots, \pi_K\}$ are referred to as the mixture proportions. N is the number of available data points and K represents the number of Gaussian distributions that are considered in the mixture. The model parameters that need to be estimated during algorithm training are

$$\boldsymbol{\theta} = \{\boldsymbol{\mu}, \boldsymbol{\Sigma}, \boldsymbol{\pi}\} \quad (2)$$

For supervised learning, each input vector, \mathbf{x} , is already labelled - in other words, the user already knows which of the Gaussians in the mixture was used to generate each sample. In such a circumstance, identifying the parameters $\boldsymbol{\theta}$ is very easy - it is shown here for illustrative purposes and to establish notation. Using X_k to denote the set of N_k samples that were generated from the k th Gaussian, the mean of each Gaussian can be estimated by

$$\boldsymbol{\mu}_k = \frac{1}{N_k} \sum_{\mathbf{x} \in X_k} \mathbf{x} \quad (3)$$

The covariance matrices are estimated using

$$\boldsymbol{\Sigma}_k = \frac{1}{N_k} \sum_{\mathbf{x} \in X_k} (\mathbf{x} - \boldsymbol{\mu}_k)(\mathbf{x} - \boldsymbol{\mu}_k)^T \quad (4)$$

while the mixture proportions are set according to

$$\pi_k = \frac{N_k}{N} \quad (5)$$

where N , as before, represents the total number data points (such that $N = \sum_k N_k$).

3.2. Unsupervised Learning

With unsupervised learning each data point is unlabelled. For a GMM model this means that, while it is assumed that each point is a sample from one of the Gaussian distributions in the mixture, the specific Gaussian distribution from which each data point was sampled is not known. In such a case the labels are described as *latent variables*, as they are hidden to the user when an analysis is conducted. This makes the problem much more difficult relative to the supervised case as now it is necessary to estimate both the parameters of the Gaussian distributions in the mixture *and* the labels associated with each data point. Difficulties arise because the parameters of the Gaussian distributions and the labels must be correlated - the geometry of the Gaussian distributions can only be estimated if the labels are known while the labels can only be estimated if the geometry of the Gaussian distributions are known.

At this point, it is convenient to write the latent variables using what is known as a 1-of-K representation. Specifically, each data point (\mathbf{x}_i , for example) is associated with a K -dimensional vector, \mathbf{z}_i . One element of \mathbf{z}_i is always equal to 1 while all the other elements of \mathbf{z}_i are set equal to zero. This means that, by stating that $z_{ik} = 1$ indicates that \mathbf{x}_i was generated from the k th Gaussian in the mixture, the set $\mathbf{Z} = \{\mathbf{z}_1, \dots, \mathbf{z}_N\}$ can be used to represent the latent variables in the problem. Further analysis can be used to show that the mixture proportions can be defined as

$$\pi_k = \Pr(z_{ik} = 1) \quad (6)$$

(see [52] for example) while the probability of observing the point \mathbf{x}_i conditional on \mathbf{z}_i and $\boldsymbol{\theta}$ is

$$p(\mathbf{x}_i | \mathbf{z}_i, \boldsymbol{\theta}) = \prod_{k=1}^K \mathcal{N}(\mathbf{x}_i; \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)^{z_{ik}} \quad (7)$$

Assuming uncorrelated samples, one can then write that

$$p(\mathbf{X} | \mathbf{Z}, \boldsymbol{\theta}) = \prod_{i=1}^N \prod_{k=1}^K \mathcal{N}(\mathbf{x}_i; \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)^{z_{ik}} \quad (8)$$

where $\mathbf{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_N\}$ is the set of all observed data. Furthermore, the posterior probability of \mathbf{Z} can be derived using Bayes' theorem:

$$p(\mathbf{Z}|\mathbf{X}, \boldsymbol{\theta}) = \frac{p(\mathbf{X}|\mathbf{Z}, \boldsymbol{\theta})p(\mathbf{Z}|\boldsymbol{\theta})}{\sum_{z \in \mathcal{Z}} p(\mathbf{X}|\mathbf{Z}, \boldsymbol{\theta})p(\mathbf{Z}|\boldsymbol{\theta})} \quad (9)$$

Equation (8) allows the maximum likelihood $\boldsymbol{\theta}$ to be identified, conditional on the latent variables. Likewise, equation (9) allows a probabilistic analysis of the latent variables, conditional on $\boldsymbol{\theta}$. This allows estimates of $\boldsymbol{\theta}$ and \mathbf{Z} to be estimated in a two-step procedure, known as the Expectation Maximization (EM) algorithm.

3.2.1. Expectation Maximization

As the name implies, the EM algorithm starts with an expectation step. Simplifying matters slightly, this is essentially where the model parameters $\boldsymbol{\theta}$ are held fixed and the expected values of the latent variables \mathbf{Z} are computed. Using equation (9) it can be shown that

$$\mathbb{E}[z_{ik}] = \frac{\mathcal{N}(\mathbf{x}_i; \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)\pi_k}{\sum_{j=1}^K \mathcal{N}(\mathbf{x}_i; \boldsymbol{\mu}_j, \boldsymbol{\Sigma}_j)\pi_j} \quad (10)$$

This step is followed by the maximization step, where the latent variables \mathbf{Z} are held equal to their expected values and the maximum likelihood of the model parameters, $\boldsymbol{\theta}$, are computed. Evaluating the derivative of equation (8) and setting the resulting expression equal to 0 then, subject to the appropriate constraints, it can be shown that the maximum likelihood parameters are

$$N_k = \sum_{i=1}^N \mathbb{E}[z_{ik}] \quad (11)$$

$$\boldsymbol{\mu}_k = \frac{\sum_{i=1}^N \mathbb{E}[z_{ik}]\mathbf{x}_i}{N_k} \quad (12)$$

$$\boldsymbol{\Sigma}_k = \frac{\sum_{i=1}^N \mathbb{E}[z_{ik}](\mathbf{x}_i - \boldsymbol{\mu}_k)(\mathbf{x}_i - \boldsymbol{\mu}_k)^T}{N_k} \quad (13)$$

$$\pi_k = \frac{N_k}{N} \quad (14)$$

The sequence of EM steps is repeated until convergence of the likelihood, equation (8), is observed. The reader may consult [52][54] for more details about the EM algorithm.

3.3. Semi-Supervised Model Formulation

In semi-supervised learning, the full data set consists of labelled and unlabelled data. The aim is to classify future data using the labelled information, while also using information contained in the unlabelled data. This approach is essentially a combination of the supervised and unsupervised formulations described in the previous sections.

For the labelled data, it is now convenient to introduce a 1-of- K representation of each label. Specifically, each labelled point \mathbf{x}_i is associated with a vector \mathbf{y}_i where, in a similar manner to our definition of the latent variables, one element of \mathbf{y}_i is always equal to 1 while all the other elements of \mathbf{y}_i are set equal to 0 (thus indicating the Gaussian that was used to generate the data point). For simplicity it is assumed that the data is ordered such that the first L points are labelled, while the remaining points are unlabelled. This allows the sets of labelled and unlabelled data to be written as

$$\{\mathbf{X}_l, \mathbf{Y}\} \equiv \{\mathbf{x}_i, \mathbf{y}_i\}_{i=1}^L \quad (15)$$

and

$$\{\mathbf{X}_u, \mathbf{Z}\} \equiv \{\mathbf{x}_i, \mathbf{z}_i\}_{i=L+1}^N \quad (16)$$

respectively. The probability of witnessing the data conditional on the GMM parameters is therefore:

$$p(\mathbf{X}_l, \mathbf{Y}, \mathbf{X}_u, \mathbf{Z} | \boldsymbol{\theta}) = p(\mathbf{X}_l, \mathbf{Y} | \boldsymbol{\theta}) p(\mathbf{X}_u, \mathbf{Z} | \boldsymbol{\theta}) \quad (17)$$

from which it is possible to show that the maximum likelihood values of $\boldsymbol{\theta}$ are

$$N_k = \underbrace{\sum_{i=1}^L y_{ik}}_{\text{Labelled data}} + \underbrace{\sum_{i=L+1}^N \mathbb{E}[z_{ik}]}_{\text{Unlabelled data}} \quad (18)$$

$$\boldsymbol{\mu}_k = \underbrace{\frac{1}{N_k} \sum_{i=1}^L y_{ik} \mathbf{x}_n}_{\text{Labelled data}} + \underbrace{\frac{1}{N_k} \sum_{i=L+1}^N \mathbb{E}[z_{ik}] \mathbf{x}_n}_{\text{Unlabelled data}} \quad (19)$$

$$\Sigma_k = \frac{1}{N_k} \underbrace{\sum_{i=1}^L y_{ik} (\mathbf{x}_i - \boldsymbol{\mu}_k) (\mathbf{x}_i - \boldsymbol{\mu}_k)^T}_{\text{Labelled data}} + \frac{1}{N_k} \underbrace{\sum_{i=L+1}^N \text{E}[z_{ik}] (\mathbf{x}_i - \boldsymbol{\mu}_k) (\mathbf{x}_i - \boldsymbol{\mu}_k)^T}_{\text{Unlabelled data}} \quad (20)$$

$$\pi_k = \frac{N_k}{N} \quad (21)$$

where underbraces have been used to highlight which parts of equations (18), (19) and (20) arise because of the labelled and unlabelled data.

The expected values of the latent variables, conditional on $\boldsymbol{\theta}$, are found using equation (10) whereby the summation is only applied to the unlabelled data. Consequently, the EM algorithm can be applied in this context (whereby the expected labels and maximum likelihood parameter estimates are updated sequentially, over a number of iterations).

4. Case Study

A Renishaw RenAM 500M L-PBF machine was used to construct two builds, each consisting of 25 individual tensile test bars. Each build involved the printing of approximately 3600 layers. Herein, these builds are referred to as B4739 and B4741 respectively.

All samples for this study were produced from a single batch of Inconel 718. Inconel 718 has a nickel mass fraction of up to 55% alloyed with iron up to 21% and chromium up to 21%. Typical properties include high strength, excellent corrosion resistance and a working temperature range between -250 °C and 650 °C. It has a wide range of applications within industry and is suitable for applications where good tensile, creep, and rupture strength is required. In particular, it is often used in situations where corrosion and oxidation resistance at high temperatures is needed. Its excellent welding characteristics and resistance to cracking makes it an ideal material for AM.

Figure 1 shows a schematic of the machine and optical system used to control the movement of the nominal 80 μm diameter focused laser spot. Samples were built in a layer-wise fashion on a substrate plate. The plate is connected to an elevator which moves vertically downwards, allowing the controlled deposition of powder layers at 60 μm intervals.

A commercially available laser processing parameter set (supplied by Renishaw) was used throughout the experiments. These were derived from standard process optimisation methods used in the AM industry. Post build, the test pieces were removed from the substrate plates using wire erosion. The tensile test bars were machined to ASTM E8-15a specification to a nominal diameter of 6.0mm and parallel length equal to 36.0mm.

Each specimen was instrumented with a dual averaging extensometer and tested at ambient temperature using an Instron tensile test machine. Tests were conducted with a 100 kN load cell under strain rate control at the first rate (0.005 strain/min) to beyond yield at which point the second rate (0.05 strain/min) was adopted, following the removal of the extensometry equipment.

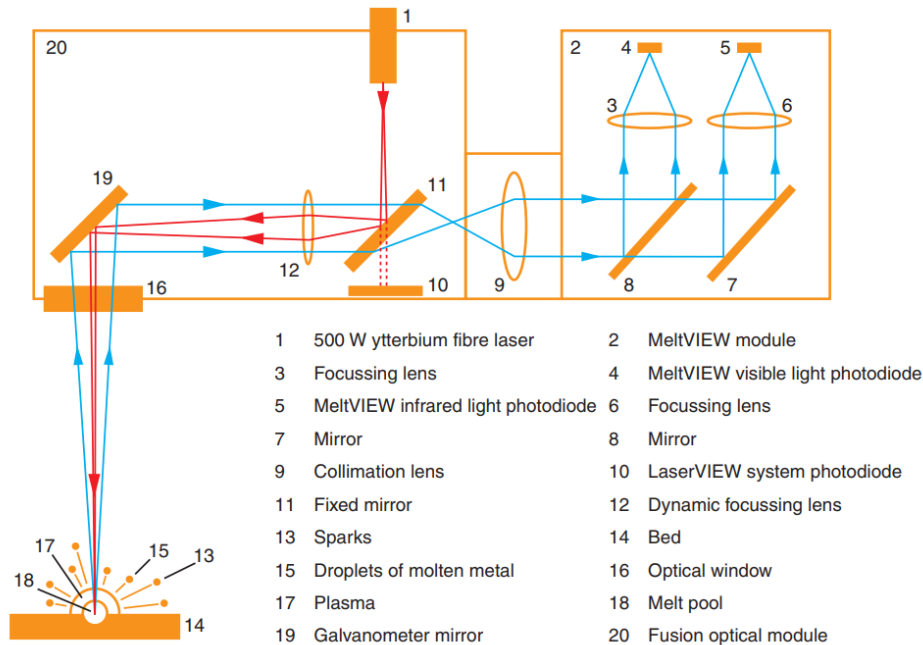


Figure 1: Schematic of Renishaw RenAM 500M L-PBF machine and optical sensing system (image reproduced with permission from the Renishaw Brochure ‘InfiniAM Spectral’, available at <http://www.renishaw.com/en/infiniam-spectral--42310>).

Figure 1 illustrates the photodiode sensing system (MeltVIEW) that was used during each build. Light from the melt pool enters the optical mirror before being reflected into the MeltVIEW module by the galvanometer mirror. A semi-transparent mirror is then used to reflect light to photodiode 1 (labelled as 4 in Figure 1) before a fully opaque mirror reflects light to photodiode 2 (labelled as 5 in Figure 1). Photodiode 1 is designed to detect plasma emissions (between 700 and 1050 nm) while photodiode 2 is designed to detect thermal radiation from the melt pool (between 1100 and 1700 nm). Time histories of the photodiode measurements and laser position were output to a series of DAT files. Each DAT file corresponded to a layer of the build and contained approximately 115 KB of data. During processing, no missing values were identified.

Using the MeltVIEW sensing system, the task is to extract significant information about build quality from the photodiode measurements. In this

work, quality is defined using the results from an Ultimate Tensile Strength (UTS) test of each bar. Here, a UTS value of 1400 MPa represents an acceptable part while UTS values below 1400 MPa represent a faulty part (this definition is sufficient for demonstrating the feasibility of the proposed approach although it is noted that more complex criteria can be utilised in the future). Figure 2 shows the x-y coordinates of the laser during the build of a single layer on the fusion bed.




Figure 2: x-y coordinates of the laser as a single layer of a build is being constructed. Red areas indicate the positions of the 25 tensile test bars while blue represents the laser path. Note that x-y coordinates are calculated from galvanometer measurements and that, for confidentiality reasons, units of position have been left as arbitrary.

Regarding the choice of sensing system, the general consensus amongst current literature is that data regarding melt pool characteristics will be closely related to build quality. Photodiode data is used in the current work as it is known to be closely correlated to properties of the melt pool (see [32], for example). While it has been hypothesised that, relative to thermal imaging systems, photodiodes may be able to capture data from a larger zone around the melt pool, the significance of this difference is currently unclear and could be investigated as future work.

4.1. Feature Extraction

As described previously, two L-PBF builds were conducted as part of the study. This led to the construction of a total of 50 tensile test bars (i.e. 25

bars per build). During each build, the x and y position of the laser was collected alongside time history measurements from 2 photodiodes sensors (sample frequency equal to 100kHz, resulting in approximately 400 GB of data per build). Here it is described how, from these large data, key features were extracted. This was based on the hypothesis that, from the photodiode measurements, it would be possible to extract relatively low dimensional features that give a statistically significant indication of build quality. It is also demonstrated how, because of the size of the data being utilised, feature extraction from L-PBF process measurements must be conducted using methods that are appropriate for large data sets.

Initial data processing / reduction was conducted in two steps. Figure 3 graphically demonstrates this process for a single build (noting that the same procedure was applied to measurements from both photodiodes, Figure 3 illustrates the process for data from a single photodiode only). In Step 1, a downsampling procedure was used such that only the data from every 10th layer of the build was used in subsequent analyses². Note that only measurements taken when the laser was active were considered. In Step 2, for each layer that was analysed, the x-y position of the laser was used to identify which parts of each photodiode measurement time history corresponded to the building of a particular tensile test bar. This data was then collected together into an $m \times n$ data matrix, \mathbf{A} , where the first column of \mathbf{A} corresponded to measurements associated with bar one, the second column of \mathbf{A} corresponded to measurements associated with bar two etc. The transpose of \mathbf{A} is illustrated graphically at the bottom of Figure 3.

²This step was taken simply to reduce the size of the data that needed to be processed. Further studies showed that using every layer did not significantly alter the results presented here.




Figure 3: Initial analysis of data from a single photodiode sensor, for a single build.

In the final step of the feature extraction procedure the intention was to apply a Singular Value Decomposition (SVD) to the data matrix, allowing \mathbf{A} to be written as the product of 3 matrices:

$$\mathbf{A} = \mathbf{U}\mathbf{D}\mathbf{V}^T \quad (22)$$

where \mathbf{U} is an $m \times n$ orthogonal matrix, \mathbf{V} is an $n \times n$ orthogonal matrix and

$$\mathbf{D} = \text{diag}(\sigma_1, \dots, \sigma_m) \quad (23)$$

where $\sigma_1, \sigma_2, \dots$ are constants (given by the eigenvalues of $\mathbf{A}^T\mathbf{A}$) that, typically, are ordered such that $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_m$. The SVD allows each of the columns in \mathbf{A} to be written as a linear combination of basis vectors. Specifically, writing $\mathbf{B} = \mathbf{D}\mathbf{V}^T$, it can be shown that

$$\mathbf{A} = \mathbf{U}\mathbf{B} \implies \mathbf{a}_j = \sum_{p=1}^n \mathbf{u}_p B_{pj} \quad (24)$$

where \mathbf{a}_j is the j th column in \mathbf{A} and \mathbf{u}_p is the p th column in \mathbf{U} . From equation (24) it can be seen that each column of \mathbf{A} is now associated with

n constants (\mathbf{a}_j is associated with $B_{1j}, B_{2j}, \dots, B_{nj}$ etc.) It is these constants that can be used as *features* - inputs to the machine learning algorithm. In fact, by ordering the SVD results such that $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_m$, close approximations of \mathbf{A} can be realised without using the full set of basis vectors. Specifically, if a new matrix, $\tilde{\mathbf{A}}$, is formed whose j th column is

$$\tilde{\mathbf{a}}_j = \sum_{p=1}^{\tilde{n}} \mathbf{u}_p B_{pj}, \quad \tilde{n} < n \quad (25)$$

then $\tilde{\mathbf{A}}$ will form a low-rank approximation of \mathbf{A} . Using $\tilde{\mathbf{A}}$ instead of \mathbf{A} can therefore facilitate a reduction in the size of the *feature space* (in other words, the number of constants associated with each column of $\tilde{\mathbf{A}}$ will be less than the number of constants associated with each column of \mathbf{A}).

Unfortunately it was found that the matrix \mathbf{A} was prohibitively large for analysis via standard SVD. To circumvent this issue \mathbf{A} was, instead, decomposed using a *Randomised SVD*. A brief outline of this procedure is given in the following text, however, for more information, readers may consult [55][56][57].

A Randomised SVD first involves the generation of an orthogonal projection matrix \mathbf{G} , which, when applied to the data matrix, reduces dimensionality while approximately preserving the pairwise distances between each of the projected vectors. To avoid the large computational costs that can be associated with this procedure, the columns of the projection matrix are sampled from a zero-mean unit-variance multivariate Gaussian distribution. This ensures that, *on average*, the required properties of the projection matrix are obtained. Once \mathbf{G} has been formed, \mathbf{A} is projected onto \mathbf{G} to realise the matrix \mathbf{H} (such that $\mathbf{H} = \mathbf{AG}$). An iterative procedure, described in [56], is then used to increase the differences between the large and small singular values of \mathbf{H} . This decreases the computational cost of the next stage of the process, whereby a QR-decomposition is used to orthonormalise the column vectors of \mathbf{H} . The QR-decomposition is used to account for the fact that, potentially, the randomly generated projection matrix \mathbf{G} may not be perfectly orthonormal. Having been orthonormalised, \mathbf{H} is then used to realise a final, low rank approximation of \mathbf{A} , denoted \mathbf{A}' . A standard SVD is then applied to \mathbf{A}' .

In the current work, for each tensile test bar, the time history of measurements from photodiodes 1 and 2 were each projected onto a *single basis vector only*. As a consequence, each specimen becomes associated with a 2-dimensional ‘feature vector’. The first element of the feature vector represents the projection of measurements from photodiode 1 onto a single basis vector while the second element of the feature vector is the projection of measurements from photodiode 2 onto a single basis vector. Inevitably, some information is lost in this projection process. Figures 4 and 5 respectively compare a segment of the measurements from photodiodes 1 and 2, for a single specimen, before and after the projection onto a single basis vector. If this level of information loss was deemed unsatisfactory one could choose to project these measurement time histories onto a greater number of basis vectors (although this would, in turn, increase the dimensionality of the space within which the machine learning algorithm must be applied). In the current study, however, it was found that projecting onto a single basis vector made it possible to distinguish between acceptable and faulty builds with sufficient accuracy - the potential benefits of projecting onto more than one basis vector will be investigated in future work.




Figure 4: Outputs for photodiode 1, for the first tensile test bar of build B4739. Black represents the uncompressed measurements, red represents measurements after they have been projected onto a single basis vector. Note that, for confidentiality reasons, unitless photodiode measurements are shown here.




Figure 5: Outputs for photodiode 2, for the first tensile test bar of build B4739. Black represents the uncompressed measurements, red represents measurements after they have been projected onto a single basis vector. Note that, for confidentiality reasons, unitless photodiode measurements are shown here.

4.2. Semi-Supervised Learning Application

Tensile tests were performed on the builds using a standard Instron tensile machine at room temperature. As detailed in Section 4.1, the ultimate tensile strength (UTS) of the bars were used to define each bar as ‘acceptable’ or ‘faulty’. Semi-supervised learning was applied to the features extracted from each of the bars. However, bar 22 from build B4741 was not considered because its ultimate tensile strength could not be obtained. As a result, 49 specimens were considered in this analysis. Figure 6 shows the position of each specimen in the feature space and the associated labels. With the aim of distinguishing between ‘acceptable’ and ‘faulty’ cases, a GMM with two Gaussian distributions was employed.

In the following, when assessing new data, specimens are labelled as faulty if the probability that they are faulty is greater than 0.5. This was considered sufficient for analysing the feasibility of the approach such that, once established, future work can aim to further exploit the uncertainty information contained in such probabilistic outputs. In the author’s opinion, it is important that an uncertainty quantification framework is built into the

proposed approach from the onset as, for approaches that are purely data based, knowing when a diagnosis is uncertain and where human intervention may be required will be crucial.




Figure 6: The position of each specimen in the feature space. The triangular points are for bars from build B4741 and the circular points are for bars from build B4739. The colour green represents acceptable specimens and red represents faulty specimens.

It is noted that, in the following, the algorithm is always initialised using the results of a purely supervised approach. Specifically, the first iteration ignores unlabelled samples and produces an initial estimate of the GMM parameters using the labelled samples only (employing equations (11), (12), (13) and (14)).

4.3. Results

Initial runs concentrated on a single case where 25 of the specimens were labelled while the remaining 24 were unlabelled. For this case, the unlabelled points were selected randomly, leading to the training data shown in Figure 7. The semi-supervised GMM was then trained, before being used to classify all 49 specimens. Using the training data shown in Figure 7, faulty specimens were identified with a 77% success rate. These results are illustrated in Figure 8, where red and green contours illustrate the positions of the two Gaussians in the mixture model, circle represents the true labels that were

assigned to each specimen and triangles show the labels inferred by the algorithm. Note that the inferred labels are colour-coded depending on the probabilities that were assigned by the algorithm - purely green triangles correspond to the probability of a faulty specimen equal to zero while purely red triangles correspond to the probability of a faulty specimen equal to one.

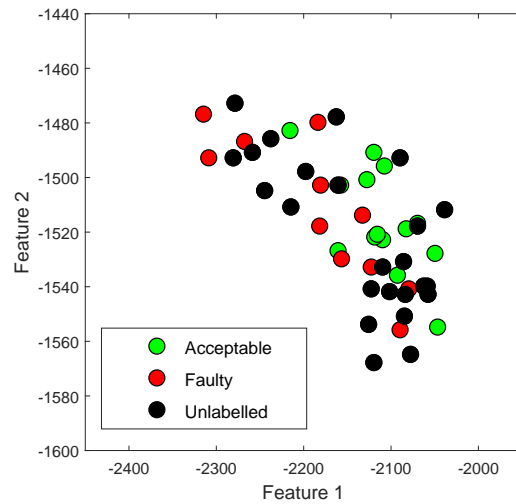


Figure 7: Training data. Example labelled and unlabelled specimens, in the feature space, before the application of semi-supervised learning.




Figure 8: Example semi-supervised learning results. Red and green contours show the inferred geometry of the two Gaussian distributions in the mixture. Circles represent the true labels that were assigned to each specimen, while triangles show the inferred labels.

The results in Figure 8 represent the algorithm's outputs for a single set of training data only. To better gauge overall performance, a Monte Carlo analysis was conducted - 1000 analyses were undertaken where, at each Monte Carlo iteration, the 24 unlabelled points were selected randomly. The resulting positions of the two Gaussian distributions were found to be relatively insensitive to the choice of unlabelled points. This is illustrated in Figure 9, which shows the results that were obtained for six runs of the Monte Carlo analysis. Furthermore, the algorithm success rate was also found to be relatively insensitive to the assignment of unlabelled data; the histogram in Figure 10 shows success rates that are closely clustered around 77%. It is important to note that, by giving a probabilistic estimate of each specimen's label, uncertainty quantification is embedded into the approach. This is useful as it can illustrate, to the user, when a particular specimen is difficult to label (i.e. when it is not clear which cluster the data point belongs to).

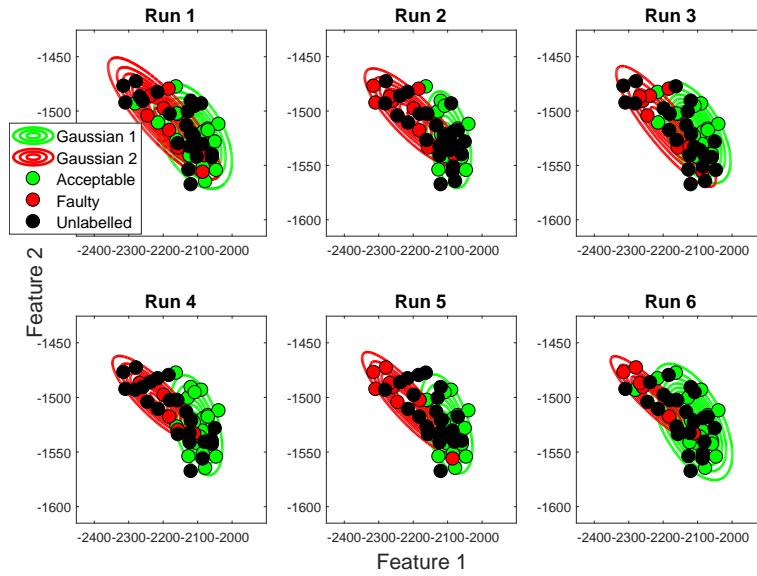


Figure 9: Semi-supervised learning results obtained for 6 runs of a Monte Carlo simulation where, for each run, 24 unlabelled points are selected randomly.




Figure 10: Histogram of algorithm success rates, obtained over 1000 runs of Monte Carlo simulation where, for each run, 24 unlabelled points are selected randomly.

To analyse how the algorithm's performance degrades as less labelled data is used, similar Monte Carlo simulations were conducted using different amounts of labelled and unlabelled data. Figure 11 shows results ranging from the case where there are 48 labelled points (and 1 unlabelled point) to the case where there are 20 label points (and 28 unlabelled points). While lower success rates are more frequently observed when the number of labelled data appoints is reduced (as one would expect), it is encouraging to note that algorithm performance does not drop off sharply. It can be seen, for example, that the number of labelled data points can be halved without significantly altering the resulting success rates. While, in the example, labels were relatively cheap to obtain (using tensile tests) the cost savings associated with the semi-supervised approach will clearly increase when more thorough and/or expensive certification methods are used. For example, in the author's experience, a CT scan of a typical component usually costs between £500 and £1000.




Figure 11: Histogram of algorithm success rates, obtained over 1000 runs of Monte Carlo simulation, as a function of the number of labelled data points.

The next set of analyses aimed to illustrate the additional information that can be acquired through the inclusion of the unlabelled data. Specifically, for the case where 16 of the data points were labelled, the success rate of the semi-supervised approach was compared with that of a purely supervised approach (where, in other words, a Gaussian Mixture Model was fitted only to the 16 labelled data). Again, a Monte Carlo analysis was used to consider variability in the selection of labelled data points. Figure 12 shows the resulting histograms of percentage success rate for (a) the semi-supervised case and (b) the supervised case. The vertical red line indicates the benchmark result i.e. the result that can be obtained if all 49 of the data points were labelled and used to train the Gaussian Mixture Model. Relative to the supervised approach, the semi-supervised approach is more clustered around the benchmark result.




Figure 12: Histogram of algorithm success rates the (a) semi-supervised, and (b) supervised only approach. Results are shown for the case where 16 of the data points were labelled. The red line indicates the benchmark result (supervised learning when all 49 points are labelled).

For each experiment used to create the results in Figure 12, a Receiver

Operating Characteristic (ROC) curve was also calculated. ROC curves are applicable when a classification algorithm outputs the *probability* of each data point belonging to a specific class (as is the case here). ROC curves involve analysing the performance of the proposed classification algorithm for a variety of ‘threshold probabilities’. In the current case, for example, success rate was defined based on a threshold probability of 0.5 as it was assumed that a specimen would be classified as ‘Faulty’ if the probability that it belonged to the corresponding ‘Faulty’ cluster exceeded 0.5. To create an ROC curve, a classifier’s False Positive Rate (FPR) and True Positive Rate (TPR) are recorded for a range of different threshold probabilities before the resulting TPR and FPR outputs are plotted against one-another. In the context of the current example, the TPR is defined as the ratio of the number of successfully identified faulty specimens, relative to the total number of identified faulty specimens. Likewise, the FPR is defined as the ratio of the number of falsely identified faulty specimens, relative to the total number of identified acceptable specimens. The closer a ROC curve is to the top-left corner of an ROC plot, the better the classifier is judged to be.

Figure 13 shows the ensemble of ROC curves that were realised for the same set of Monte Carlo experiments used to create Figure 12 (where there were 16 labelled points and 33 unlabelled points). To aid visualisation, the individual ROC curve evaluations are plotted using markers, where blue stars represent results obtained using a semi-supervised approach and black circles represent results obtained using a supervised approach, applied to the 16 labelled data points. The red ROC curve represents the baseline case, where all 49 of the data points were labelled and used to train the Gaussian Mixture Model. Similarly to the results shown in Figure 12, the ensemble of results using the semi-supervised approach cluster the benchmark result more closely than those realised using the supervised approach.

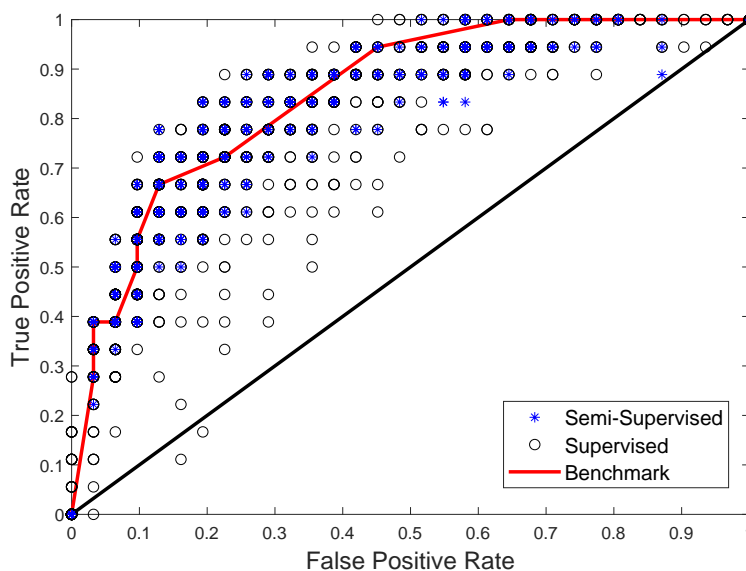


Figure 13: ROC curves for the ensemble of experiments analysed in Figure 12. The red line indicates the benchmark result (supervised learning when all 49 points are labelled).

To complete a comparison of the supervised and semi-supervised approaches, 2-fold cross validation was performed. In other words, the data was divided into two ‘training’ and ‘testing’ data sets of 25 and 24 points respectively. This data separation allowed us to analyse the ability of the algorithm to classify data in the testing set, based only on information from the training set. To complete the 2-fold validation, the role of the two datasets was reversed (such that the training data became the testing data and so on), such that a second analysis could be performed. The success rates for these two analyses were then averaged, to obtain a measure of algorithm performance.

As one would expect, the differences between the semi-supervised and supervised approaches become more pronounced as the ratio of unlabelled training data to labelled training data increases. For this reason, for illustrative purposes, Figure 14 represents results obtained for the case where only 6 of the training data points were labelled. As before, to account for variations that may arise in the selection of labelled training points, Monte

Carlo experiments were performed where, at each iteration, the labelled data was randomly selected from the available training data.

Figure 14 (a) and (b) shows the ensemble of results that were obtained for the semi-supervised and supervised cases respectively. Once again, the semi-supervised results are clustered towards a higher average success rate than the supervised results.

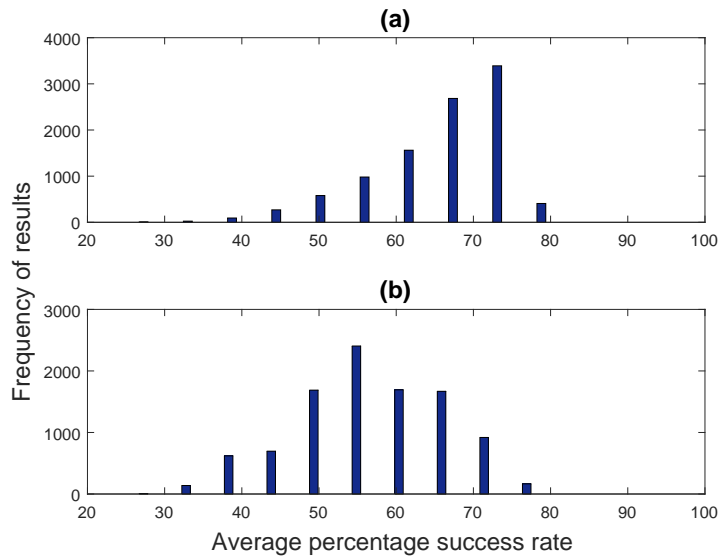


Figure 14: 2-fold cross validation where (a) are results realised using a semi-supervised approach and (b) are results using a supervised approach.

5. Conclusion and Future Work

Additive Manufacturing (AM) is a digital approach for manufacturing highly customised components. However, uncertainties surrounding part quality hinders the adoption of AM technology in many risk-averse sectors. This paper is the outcome of a feasibility study wherein a semi-supervised machine learning algorithm was developed and applied to a large amount of AM process data (photodiode measurements, generated during L-PBF

builds of tensile test bars). Key features were extracted from these large datasets using a Randomised Singular Value Decomposition, before a Gaussian Mixture Model was trained to recognise builds that had been identified as ‘faulty’. The semi-supervised approach allowed this to be conducted using a reduced number of certification experiments and, through Receiver Operating Characteristic (ROC) curves and 2-fold cross-validation, was shown to outperform a supervised approach that did not use the unlabelled data that was available for training. Key contributions are summarised as follows:

1. In this work it was demonstrated how, when using machine learning to infer part quality from L-PBF process measurements, the large quantity of available data can prevent the application of ‘conventional’ feature extraction methods. It was illustrated how this challenge can be overcome using methods that are suitable for large datasets (a Randomised Singular Value Decomposition in this case).
2. By successfully classifying ‘successful’ builds with a 77% success rate, the feasibility of identifying faulty L-PBF builds using a purely data-based approach analysis of photodiode measurement time histories has been demonstrated.
3. It has been demonstrated that, through a semi-supervised approach, the number of costly certification experiments required in the implementation of machine-learnt build classification can be significantly reduced.

The paper has led to several avenues of future work.

Firstly, the authors are currently investigating whether the results reported in the current manuscript can be improved through the use of additional basis vectors. This will reduce the amount of information lost during feature extraction but will also increase the dimensionality of the feature space within which machine learning must be performed. Secondly, with regard to sensing systems, the current paper utilised data from photodiodes sensors (which has been shown to be closely related properties of the melt pool [32]). Future work aims to investigate whether classification can be improved through the use of additional, complimentary sensing systems (acoustic sensors and thermal imaging cameras, for example). Finally, the

authors are currently developing a version of the semi-supervised algorithm described in the current paper that is suitable for *layer-by-layer* defect detection, using data provided from CT scans. Ultimately, the aim of this work is to establish machine-learned control strategies that can de-risk AM Technology, facilitate its wider adoption and reduce the time associated with new materials innovation.

Acknowledgement

This work was conducted in the feasibility study ‘Towards Additive Manufacturing Process Control using Semi-Supervised Machine Learning’ that was funded through the EPSRC Network Plus Grant: Industrial Systems in the Digital Age EP/P001246/1, <https://connectedeverything.ac.uk/>.

References

- [1] E. Uhlmann, R. Pontes, A. Laghmouchi, A. Bergmann, Intelligent pattern recognition of a SLM machine process and sensor data, *Procedia CIRP* 62 (2017) 464–469.
- [2] P. O’Regan, P. Prickett, R. Setchi, G. Hankins, N. Jones, Metal based additive layer manufacturing: variations, correlations and process control, *Procedia Computer Science* 96 (2016) 216–224.
- [3] R. Rashid, S. Masood, D. Ruan, S. Palanisamy, R. Rashid, M. Brandt, Effect of scan strategy on density and metallurgical properties of 17-4ph parts printed by Selective Laser Melting (SLM), *Journal of Materials Processing Technology* 249 (2017) 502–511.
- [4] V. Renken, S. Albinger, G. Goch, A. Neef, C. Emmelmann, Development of an adaptive, self-learning control concept for an additive manufacturing process, *CIRP Journal of Manufacturing Science and Technology* 19 (2017) 57–61.
- [5] D. Lieber, M. Stolpe, B. Konrad, J. Deuse, K. Morik, Quality prediction in interlinked manufacturing processes based on supervised & unsupervised machine learning, *Procedia Cirp* 7 (2013) 193–198.

- [6] N. Sadati, R. Chinnam, M. Nezhad, Observational data-driven modeling and optimization of manufacturing processes, *Expert Systems with Applications* 93 (2018) 456–464.
- [7] J. Pang, Y. Gu, J. Xu, G. Yu, Semi-supervised multi-graph classification using optimal feature selection and extreme learning machine, *Neurocomputing* 277 (2018) 89–100.
- [8] Y. Chen, B. Zhao, J. Zhang, Y. Zheng, Automatic segmentation for brain mr images via a convex optimized segmentation and bias field correction coupled model, *Magnetic resonance imaging* 32 (7) (2014) 941–955.
- [9] Z. Ji, Y. Huang, Y. Xia, Y. Zheng, A robust modified gaussian mixture model with rough set for image segmentation, *Neurocomputing* 266 (2017) 550–565.
- [10] B. Ribeiro, Support vector machines for quality monitoring in a plastic injection molding process, *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)* 35 (3) (2005) 401–410.
- [11] X. Liu, Q. Xu, J. Ma, H. Jin, Y. Zhang, Mslrr: a unified multiscale low-rank representation for image segmentation, *IEEE Transactions on Image Processing* 23 (5) (2014) 2159–2167.
- [12] D. Montgomery, *Statistical quality control*, Vol. 7, Wiley New York, 2009.
- [13] W. Sames, F. List, S. Pannala, R. Dehoff, S. Babu, The metallurgy and processing science of metal additive manufacturing, *International Materials Reviews* 61 (5) (2016) 315–360.
- [14] T. Spears, S. Gold, In-process sensing in selective laser melting (SLM) additive manufacturing, *Integrating Materials and Manufacturing Innovation* 5 (1) (2016) 2.
- [15] T. DebRoy, H. Wei, J. Zuback, T. Mukherjee, J. Elmer, J. Milewski, A. Beese, A. Wilson-Heid, A. De, W. Zhang, Additive manufacturing of metallic components—process, structure and properties, *Progress in Materials Science*.

- [16] G. Casalino, Computational intelligence for smart laser materials processing, *Optics & Laser Technology* 100 (2018) 165–175.
- [17] S. Negi, S. Dhiman, R. Sharma, Determining the effect of sintering conditions on mechanical properties of laser sintered glass filled polyamide parts using rsm, *Measurement* 68 (2015) 205–218.
- [18] R. Enneti, R. Morgan, S. Atre, Effect of process parameters on the selective laser melting (SLM) of tungsten, *International Journal of Refractory Metals and Hard Materials* 71 (2018) 315–319.
- [19] Y. Tian, D. Tomus, P. Rometsch, X. Wu, Influences of processing parameters on surface roughness of hastelloy x produced by selective laser melting, *Additive Manufacturing* 13 (2017) 103–112.
- [20] J. Tan, W. Wong, K. Dalgarno, An overview of powder granulometry on feedstock and part performance in the selective laser melting process, *Additive Manufacturing*.
- [21] E. Olakanmi, R. Cochrane, K. Dalgarno, A review on selective laser sintering/melting (SLS / SLM) of aluminium alloy powders: Processing, microstructure, and properties, *Progress in Materials Science* 74 (2015) 401–477.
- [22] X. Liu, M. Zeng, Y. Ma, M. Zhu, Wear behavior of al–sn alloys with different distribution of sn dispersoids manipulated by mechanical alloying and sintering, *wear* 265 (11-12) (2008) 1857–1863.
- [23] F. Caltanissetta, M. Grasso, S. Petr, B. M. Colosimo, Characterization of in-situ measurements based on layerwise imaging in laser powder bed fusion, *Additive Manufacturing* 24 (2018) 183 – 199. doi:<https://doi.org/10.1016/j.addma.2018.09.017>.
URL <http://www.sciencedirect.com/science/article/pii/S2214860418305785>
- [24] B. Zhang, J. Ziegert, F. Farahi, A. Davies, In situ surface topography of laser powder bed fusion using fringe projection, *Additive Manufacturing* 12 (2016) 100 – 107. doi:<https://doi.org/10.1016/j.addma.2016.08.001>.
URL <http://www.sciencedirect.com/science/article/pii/S2214860416300720>

- [25] Z. Li, X. Liu, S. Wen, P. He, K. Zhong, Q. Wei, Y. Shi, S. Liu, In situ 3d monitoring of geometric signatures in the powder-bed-fusion additive manufacturing process via vision sensing methods, *Sensors* 18 (4). doi: 10.3390/s18041180.
URL <http://www.mdpi.com/1424-8220/18/4/1180>
- [26] S. Clijsters, T. Craeghs, S. Buls, K. Kempen, J. Kruth, In situ quality control of the selective laser melting process using a high-speed, real-time melt pool monitoring system, *The International Journal of Advanced Manufacturing Technology* 75 (5-8) (2014) 1089–1101.
- [27] U. Thombansen, A. Gatej, M. Pereira, Process observation in fiber laser-based selective laser melting, *Optical Engineering* 54 (1) (2014) 011008.
- [28] I. Yadroitsev, P. Krakhmalev, I. Yadroitsava, Selective laser melting of ti6al4v alloy for biomedical applications: Temperature monitoring and microstructural evolution, *Journal of Alloys and Compounds* 583 (2014) 404 – 409. doi:<https://doi.org/10.1016/j.jallcom.2013.08.183>.
URL <http://www.sciencedirect.com/science/article/pii/S0925838813020689>
- [29] J. A. Kanko, A. P. Sibley, J. M. Fraser, In situ morphology-based defect detection of selective laser melting through inline coherent imaging, *Journal of Materials Processing Technology* 231 (2016) 488 – 500. doi:<https://doi.org/10.1016/j.jmatprotec.2015.12.024>.
URL <http://www.sciencedirect.com/science/article/pii/S0924013615302387>
- [30] M. Grasso, B. M. Colosimo, Process defects and in situ monitoring methods in metal powder bed fusion: a review, *Measurement Science and Technology* 28 (4). doi:10.1088/1361-6501/aa5c4f.
- [31] T. Craeghs, F. Bechmann, S. Berumen, J. Kruth, Feedback control of layerwise laser melting using optical sensors, *Physics Procedia* 5 (2010) 505–514.
- [32] J. Kruth, P. Mercelis, J. Van Vaerenbergh, T. Craeghs, Feedback control of selective laser melting, in: *Proceedings of the 3rd international conference on advanced research in virtual and rapid prototyping, 2007*, pp. 521–527.

- [33] M. Mahesh, Y. Wong, J. Fuh, H. Loh, A six-sigma approach for benchmarking of rp&m processes, *The International Journal of Advanced Manufacturing Technology* 31 (3-4) (2006) 374–387.
- [34] Y. Ning, Y. Wong, J. Fuh, H. Loh, An approach to minimize build errors in direct metal laser sintering, *IEEE Transactions on automation science and engineering* 3 (1) (2006) 73–80.
- [35] C. Zhang, L. Li, A coupled thermal-mechanical analysis of ultrasonic bonding mechanism, *Metallurgical and Materials Transactions B* 40 (2) (2009) 196–207.
- [36] C. Ning, F. You, Data-driven decision making under uncertainty integrating robust optimization with principal component analysis and kernel smoothing methods, *Computers & Chemical Engineering*.
- [37] A. Karimi, C. Kammer, A data-driven approach to robust control of multivariable systems by convex optimization, *Automatica* 85 (2017) 227–233.
- [38] F. Tao, J. Cheng, Q. Qi, M. Zhang, H. Zhang, F. Sui, Digital twin-driven product design, manufacturing and service with big data, *The International Journal of Advanced Manufacturing Technology* (2017) 1–14.
- [39] Z. Ge, Review on data-driven modeling and monitoring for plant-wide industrial processes, *Chemometrics and Intelligent Laboratory Systems*.
- [40] J. Wang, Y. Ma, L. Zhang, R. Gao, D. Wu, Deep learning for smart manufacturing: Methods and applications, *Journal of Manufacturing Systems*.
- [41] M. Sharp, R. Ak, T. Hedberg, A survey of the advancing use and development of machine learning in smart manufacturing, *Journal of Manufacturing Systems*.
- [42] M. Grasso, V. Laguzza, Q. Semeraro, B. Colosimo, In-process monitoring of selective laser melting: spatial detection of defects via image data analysis, *Journal of Manufacturing Science and Engineering* 139 (5) (2017) 051001.

- [43] M. Khanzadeh, S. Chowdhury, M. Tschopp, H. Doude, M. Marufuzzaman, L. Bian, In-situ monitoring of melt pool images for porosity prediction in directed energy deposition processes, *IISE Transactions* (2018) 1–19.
- [44] L. Scime, J. Beuth, Anomaly detection and classification in a laser powder bed additive manufacturing process using a trained computer vision algorithm, *Additive Manufacturing* 19 (2018) 114–126.
- [45] M. Khanzadeh, S. Chowdhury, M. Marufuzzaman, M. Tschopp, L. Bian, Porosity prediction: Supervised-learning of thermal history for direct laser deposition, *Journal of manufacturing systems* 47 (2018) 69–82.
- [46] G. Tapia, A. Elwany, H. Sang, Prediction of porosity in metal-based additive manufacturing using spatial gaussian process models, *Additive Manufacturing* 12 (2016) 282–290.
- [47] C. Gobert, E. Reutzel, J. Petrich, A. Nassar, S. Phoha, Application of supervised machine learning for defect detection during metallic powder bed fusion additive manufacturing using high resolution imaging., *Additive Manufacturing* 21 (2018) 517–528.
- [48] S. Shevchik, C. Kenel, C. Leinenbach, K. Wasmer, Acoustic emission for in situ quality monitoring in additive manufacturing using spectral convolutional neural networks, *Additive Manufacturing* 21 (2018) 598–604.
- [49] D. Ye, G. S. Hong, Y. Zhang, K. Zhu, J. Y. H. Fuh, Defect detection in selective laser melting technology by acoustic signals with deep belief networks, *The International Journal of Advanced Manufacturing Technology* (2018) 1–11.
- [50] G. McLachlan, T. Krishnan, *The EM algorithm and extensions*, Vol. 382, John Wiley & Sons, 2007.
- [51] J. Yu, C. Chaomurilige, M. Yang, On convergence and parameter selection of the em and da-em algorithms for gaussian mixtures, *Pattern Recognition* 77 (2018) 188–203.
- [52] C. Bishop, *Pattern Recognition and Machine Learning.*, Springer-Verlag New York, 2016.

- [53] X. Hong, B. Huang, Y. Ding, F. Guo, L. Chen, L. Ren, Multi-model multivariate gaussian process modelling with correlated noises, *Journal of Process Control* 58 (2017) 11–22.
- [54] J. Ma, J. Chen, W. Xiong, F. Ding, Expectation maximization estimation algorithm for hammerstein models with non-gaussian noise and random time delay from dual-rate sampled-data, *Digital Signal Processing* 73 (2018) 135–144.
- [55] D. Achlioptas, Database-friendly random projections: Johnson-lindenstrauss with binary coins, *Journal of computer and System Sciences* 66 (4) (2003) 671–687.
- [56] N. Halko, P. Martinsson, Y. Shkolnisky, M. Tygert, An algorithm for the principal component analysis of large data sets, *SIAM Journal on Scientific computing* 33 (5) (2011) 2580–2594.
- [57] P. Indyk, R. Motwani, Approximate nearest neighbors: towards removing the curse of dimensionality, in: *Proceedings of the thirtieth annual ACM symposium on Theory of computing*, ACM, 1998, pp. 604–613.