Rapid Prototyping J

An alloy-agnostic machine learning framework for process mapping in laser powder bed fusion

Journal:	*Rapid Prototyping Journal*
Manuscript ID	RPJ-02-2024-0068.R1
Manuscript Type:	Original Article
Keywords:	Machine Learning, Additive Manufacturing, Laser Powder Bed Fusion, Process Optimisation

Note: The following files were submitted by the author for peer review, but cannot be converted to PDF. You must view these files (e.g. movies) online.

single_tracks.tex

An alloy-agnostic machine learning framework for process mapping in laser powder bed fusion

Abstract

Purpose - This study introduces an image-based method to determine the processing window for a given alloy system using laser powder bed fusion equipment based on achieving the desired melting mode across multiple materials for powder-free specimens. The method uses a convolutional neural network trained to classify different track morphologies across different alloy systems to select appropriate printing settings. This method is intended for the development of new alloy systems, where the powder feedstock may be unavailable, or prohibitively expensive to manufacture.

Design/methodology/approach - A convolutional neural network is designed from scratch to identify the 4 key melting modes that are observed in laser powder bed fusion additive manufacturing across different alloy systems. To increase the prediction accuracy and generalisation accuracy across different materials, the network is trained using a novel hybrid dataset that combines fully unsupervised learning with semi-supervised learning.

Findings - This study demonstrates that our convolutional network with a novel hybrid training approach can be generalised across different materials, and k-fold validation shows that the model retains good accuracy with changing training conditions. The model can predict the processing maps for the different alloys with an accuracy of up to 96% in some cases. It is also shown that powder-free single-track experiments are a useful indicator for predicting the final print quality of a component.

Originality - The "invariant information clustering" (IIC) approach is applied to process optimisation for additive manufacturing. A novel hybrid dataset construction approach that accounts for uncertainty in the ground truth data, enables the trained convolutional model to perform across a range of different materials, and most importantly, generalise to materials outside of the training dataset. Compared to the traditional cross-sectioning approach, this method considers the whole length of the single track when determining the melting mode.

Keywords: Laser Powder Bed Fusion, Additive Manufacturing, Machine Learning, Process Optimisation

1. Introduction

Laser Powder Bed Fusion (LPBF) is a type of additive manufacturing (AM) that uses a laser beam to melt regions of a metal powder feedstock into a solid object [1]. In direct contrast to traditional manufacturing methods, the solid material is placed only where required during the manufacture, either to form the component or to provide a support structure. This allows much greater geometric freedom in design, permitting engineers to use advanced design techniques, such as topology optimisation [2–4] and lattice structures [5, 6].

However, component design is only part of this process. To fabricate high-quality, reproducible components, the manufacturing route must first be optimised through careful selection of the machine settings. Researchers usually accomplish this parameter optimisation by making small test samples, such as cubes [7–10] or tensile bars [11, 12], to identify the parameters that can result in high-density components. Printing tensile samples is an important step in qualifying an alloy or process in specific industrial settings [13] and calibrating computational models. However, fabricating 3D samples can present challenges, such as build failures, risks associated with handling powdered metals, and the requirement of a minimum volume of powder alloy to complete the optimisation process. This is particularly problematic for new alloy development, where the minimum powder volume required to fill the LPBF machine may not be available, or it is not economically viable to produce. Therefore, this paper aims to develop a machine-learning-based framework to obtain process maps from powder-free single scan tracks using solid alloy samples, avoiding the risks and costs associated with fabricating 3D samples and metallic powders.

Machine learning-assisted process optimization is an active research topic in the AM community. Researchers have used unsupervised machine learning (ML) models to cluster images of thin walls printed with copper and to determine the optimal parameter set for the copper in the work of Silbernagel et al. [14]. Gaussian Process Regression (GPR), a statistical regression method, has also been used to create process maps based on the measured density of printed cubes against the laser power and scan speed settings [15, 16]. Some studies have extended this approach by fabricating tensile bars and measuring the variation in tensile strength of the samples with the processing parameters to refine further the parameter selection, such as the work of Yakout et al. [17].

Image-based ML has been used extensively for process monitoring and defect detection within AM, such as in the work of Scime and Beuth [18].

In addition to ML-assisted process optimisation, another method investigates the morphology of the melt tracks left behind by the laser during the printing process and theorises that the ideal melting morphology will result in the highest-quality parts. Certain types of porosity and print defects observed in LPBF directly result from undesirable melting modes caused by improper printing parameters, and these melting modes are visible in the cross-section of the melt tracks. Scime and Beuth [19] have multiple publications in this field, and detail these different cross-sectional morphologies for the Inconel 718 alloy. In the work of Gong et al. [20], the authors follow a procedure to print single tracks with and without powder on substrate plates to develop a parameter optimisation process. Their work concludes that the melting mode for tracks deeper than the 30μm powder layer is consistent with or without powder, but the powder's presence impacts the surface roughness. Importantly, to determine the melting mode of the tracks, the authors needed to cut the samples and prepare the cross-section to view the shape of the melt pool, which is labour-intensive.

Figure 1 shows the typical process map and melting modes for alloys processed using LPBF. In the literature, 4 zones are typically observed:

- **Lack-of-Fusion**: insufficient energy from the laser results in poor inter-layer bonding, and porosity.
- **Keyhole**: excessive energy causes the vaporisation of the alloy and high penetration into the substrate. The vaporised alloy is then trapped when the surrounding alloy solidifies, creating a pore.
- **Balling**: high scanning velocity creates an elongated melt pool, and then surface tension causes this to break up into multiple smaller pools [11].
- **Optimal**: in this zone, the melt pool is relatively stable, and there are minimal (if any) of the defects listed above. This is where high part density is usually found.

The exact locations of these zones will change depending on the alloy properties and the equipment being used, but generally, they will be located as shown in Figure 1. Importantly, methods to determine these melting modes rely on a cross-sectional view of the melt track, as detailed in the work of Scime et al. in [19].

Where these studies fall short is the application of knowledge learned from one material to the prediction of the behaviour of other materials. Typically, a neural network or Gaussian model will need to be retrained per material, and will only be useful with that one material. The authors of this paper believe that there is a more elegant solution, enabling engineers to enrich a single model with knowledge gained from multiple alloy systems to predict the behaviour of new material formulations.

This work proposes an unsupervised ML framework to determine the optimal processing window for a new alloy based on top-down images of single tracks printed onto a solid substrate. This method does not use powder, meaning there is no dangerous material to handle, the experimental set-up is much faster, and expensive powder atomisation is not required. For new alloys in development, a suitable powder feedstock may not be available, or in sufficient quantities to enable a comprehensive parameter optimisation process. This framework removes that requirement and uses only solid samples. Scanning individual tracks also takes a fraction of the time to execute compared to solid samples, and comes with an almost negligible risk of build failure, since no parts are being produced. The impact of neglecting the powder from the single track study is discussed in more detail in Section 4.3.

Secondly the proposed convolutional network accepts sections of the top-down view of the single tracks as input. This means that more of the track can be considered when determining the melting mode. The traditional method of taking a cross-section and measuring the width and depth only considers a single cross-section of the track. This

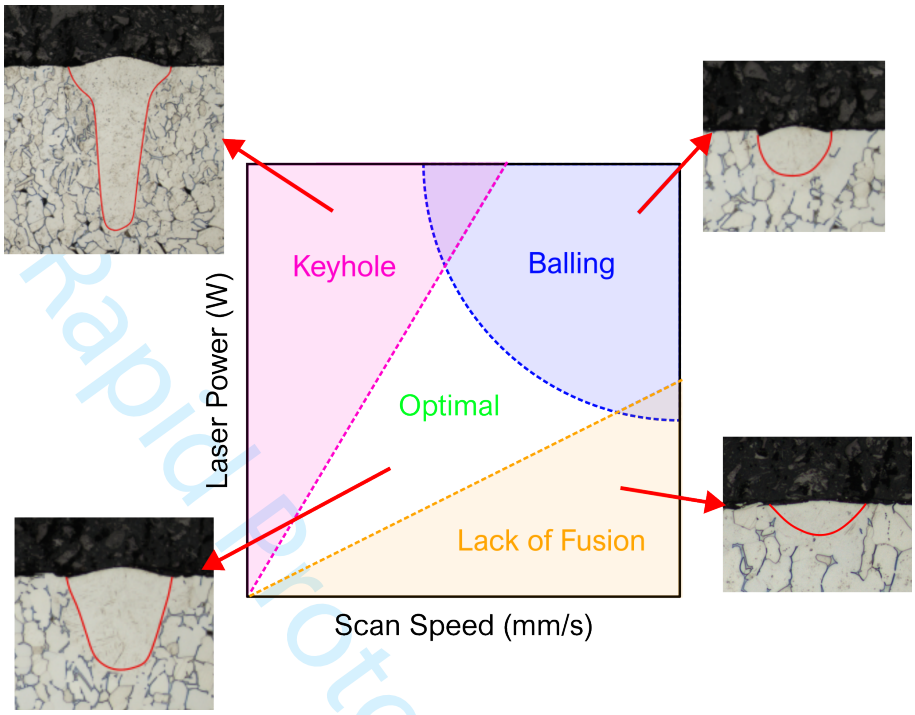

Figure 1: The typical processing map for metals processed using LPBF with the usual melt track morphologies for Ti-6Al-4V indicated. Straight dashed lines on this plot have constant energy density (except the blue balling boundary). *Authors own work.*

cross-section may vary across the length of the track. Taking multiple cross-sections is a very time-intensive process. Instead, the whole track length can be accounted for when determining the melting mode by using sections from the top-down view. The only alternative approach to determining the section across the entire track length would be to use a high-resolution CT scanner, which again is time-intensive, but also requires the hardware and significant peripheral equipment.

The paper is structured into 3 sections; Section 2 discusses the alloy selection, the parameter selection for the single tracks, and the data acquisition process; Section 3 introduces the ML framework and the image processing, as well as how the results from the neural network are interpreted; and finally, Section 4 discusses the generated process maps and how the performance of the model varies in different training cases. Section 4 also discusses the application of this method to 3D samples, and how well the knowledge learned from single-track studies translates to predicting the final part quality.

2. Experimental methodology

In this section, all the relevant details regarding the experimental part of the work will be discussed. This includes the chosen alloys, parameter selection, scanning the single tracks, and finally image acquisition.

2.1. Printing

The single-track samples in this study were all produced using an Aconity Mini (Aconity GmBH, Germany) LPBF system. This machine has a 400W IR single-mode laser with a Gaussian beam. The build chamber is purged with argon gas to an oxygen concentration below 200ppm and maintained there for each experiment. No preheating of the substrate was used. The surface of each substrate was ground to a finish of 100 grit before the tracks were printed and each sample was aligned to the build chamber before the laser was activated to ensure a consistent 80μm spot diameter.

Alloy	Conductivity (W/m·K)
Al-Si-10Mg	150 [9]
β-Ti	Similar to Ti-6Al-4V
SS316L	16.2 [23]
Ti-6Al-4V	7.07 [24]

Table 1: Thermal conductivity values for the alloys used in the present study. The value for the β-titanium alloy is not found in the literature, so it is assumed to be similar to the Ti-6Al-4V. *Authors own work.*

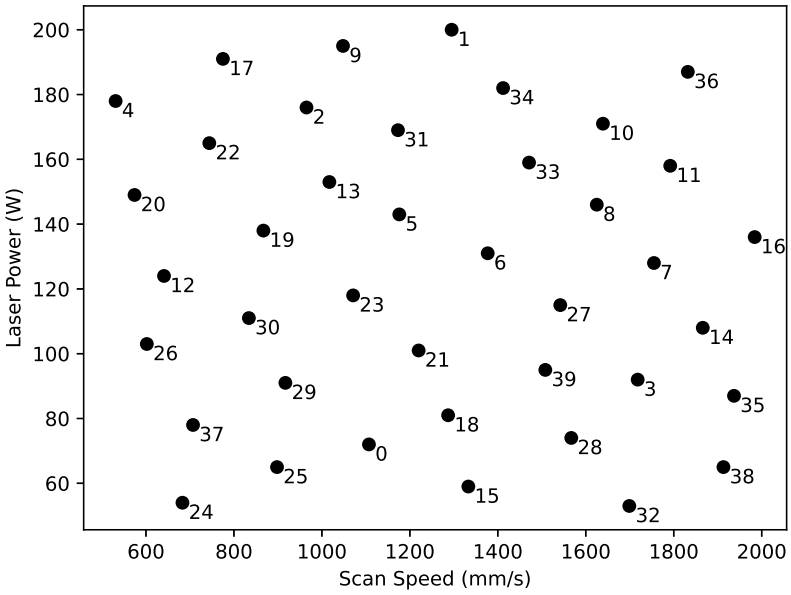

Figure 2: The design of experiments for the Ti-6Al-4V alloy single tracks. *Authors own work.*

The alloys used in the study are Ti-6Al-4V, Al-Si-10Mg, 316L stainless steel and a β-stable Ti Alloy [21]. These were chosen since they represent some of the most commonly seen alloys in the metal LPBF community. The thermal properties of these alloys will influence the behaviour of each during the printing process. Table 1 contains the thermal conductivity values in the literature for these alloys. The conductivity of the β-titanium alloy could not be found in the literature, and its measurement is not in the scope of this work. We assume its value is close to the Ti-6Al-4V and will have similar melting behaviour. Trapp et al. [22] discuss how much lower the laser absorptivity is for Al-Si-10Mg when compared to 316L stainless steel and a tungsten alloy in their work. This thermal parameter will affect the melting behaviour and the resulting images from this study.

As mentioned in the introduction, only 2 parameters are varied in this study: the laser power (P) and the scanning speed (v). The spot diameter was fixed at 80µm for all of the experiments. The parameter sets used to process each alloy are included in Appendix Appendix B. For each alloy, a new set of sample points is selected in the parameter space using an optimised Latin-hypercube sampling scheme from the Surrogate Modelling Toolbox by Bouhlel et al. [25]. This sampling scheme was optimised using the Enhanced Stochastic Evolutionary algorithm (ESE) developed by Jin et al. [26] to provide an even distribution of points in the space. The parameter set for the Ti-6Al-4V alloy single tracks is plotted in Figure 2.

Due to the high thermal conductivity of the Al-Si-10Mg relative to the other alloys, the parameter space of this alloy has been expanded to ensure that all 4 melt morphologies are observed in the track images.

The build files for each alloy were created using a Python script. The tracks were spaced 0.35mm apart so that

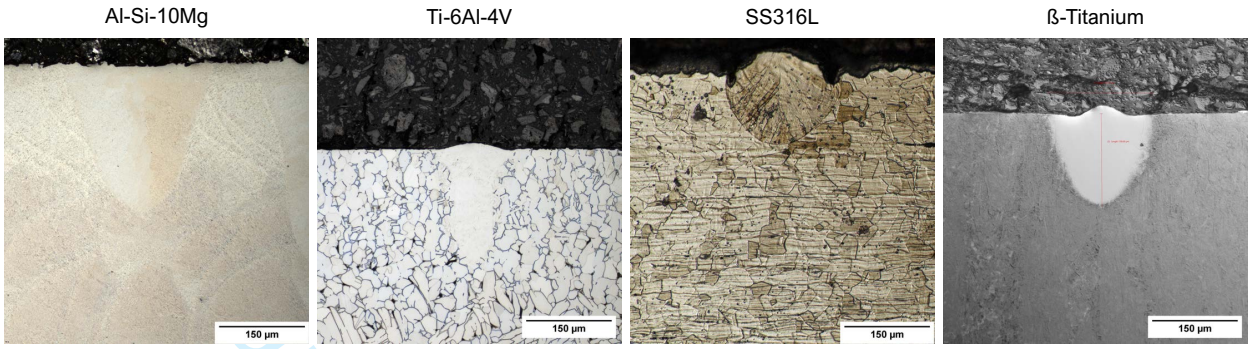

Figure 3: Cross-sections of each of the 4 chosen alloys showing the melt pool. *Authors own work.*

multiple could be imaged simultaneously with the optical microscope. They are all 10mm in length.

2.2. *Ground truth labelling*

Knowledge of the true melting mode of each track in this work is imperative to understanding the performance of any predictive workflow. The standard procedure to determine the melting mode of a single track is to obtain the cross-section of the melt pool and calculate its aspect ratio [27–29]. The samples were cut with a metallurgical saw, ground and polished up to colloidal silica before being etched with standard solutions for each alloy so the microstructure and melt track were visible, as shown in Figure 3. Using ImageJ software [30], the width w and depth d of the melt pool for each track were measured. The aspect ratio, $R = d/w$, between the depth and width was calculated for each image.

Patel et al. [29] develop a classification for single tracks printed using LPBF equipment by defining a transition region between conduction and keyhole mode from an aspect ratio $0.5 < R < 0.8$, with full conduction mode melting occurring when $R < 0.5$ and full keyhole melting occurring when $R > 0.8$. In the same work, they discuss that the final part is likely to contain lack-of-fusion defects when in conduction mode, so this region is labelled as undermelting. Based on these studies, this work uses the following thresholds to define the melting mode: undermelting where $R < 0.5$; optimal when $0.5 < R < 0.8$; keyhole when $R > 0.8$. Balling tracks are identified by their characteristic shape, with the centre of the melt track protruding above the substrate and depressed areas on either side of this protrusion.

A single cross-section image was taken for each track, approximately in the centre of the length to avoid the laser on/off affecting the mode. As mentioned in the introduction, this single cross-section image may not represent the entire track, and the dimensions may change along the length due to defects in the substrate, the surface texture, and other stochastic fluid effects occurring in the melt pool zone. Hence, there is a level of uncertainty in the class assigned to each track. This is taken into account by using a soft clustering approach, which is discussed in more detail in Section 3.4.

2.3. *Image acquisition and processing*

One of the primary difficulties in applying machine learning methods in an additive manufacturing context is the availability of a large enough dataset to train a model with many parameters [31]. This paper addresses this problem by taking advantage of the fact that tracks vary along the length, and by slicing them along their lengths, many more image samples can be obtained.

After fabrication, the samples were imaged using an optical microscope (Olympus GX53 in bright field mode). As mentioned before, the placement of the tracks allowed 3 to be imaged simultaneously, speeding up the acquisition process, as shown in Figure 4A. Approximately 16 images were captured for each track along the length, including the ends. All images were captured at a resolution of 2168x2168 pixels.

The first step in the dataset preparation is to remove the background from the images, leaving only the single tracks. This step is performed manually using ImageJ [30], as shown in Figure 4B & C. Then, the images are reduced in size to 25% of their original dimensions to preserve GPU memory and allow larger batch sizes for more parallel training. Next, the image is sliced into 3 sections to separate the 3 tracks, and then each track is cut into 20 slices along

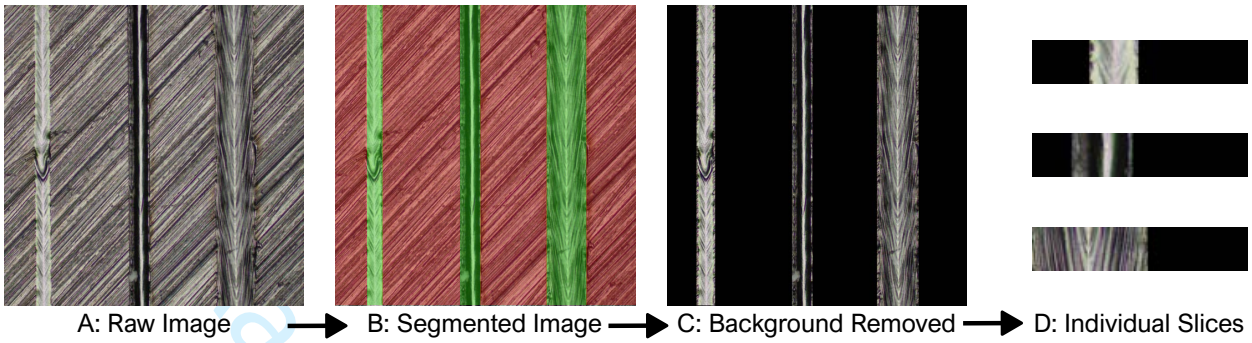

Figure 4: An overview of the image processing workflow used in this work. The raw input image (A) is manually segmented (B) to remove the background (C) before being sliced into individual samples (D). *Authors own work.*

the length. This results in approximately 320 individual samples per track or around 13,000 images per alloy. These smaller images are further cropped to remove more of the background, and finally measure 136x27 pixels, shown in Figure 4D. The images are saved to disk as 8-bit ".png" files.

3. Computational framework

This paper proposes a deep neural network to accomplish the unsupervised clustering of the track images. Convolutional neural networks (CNNs) have been shown to perform well at various tasks involving images [32, 33], due to their ability to perform feature engineering automatically, selecting the most important parts of the image that contribute to a final classification. This feature engineering also happens across multiple length scales, since CNNs use sub-sampling to reduce the size of the image with successive layers, either through variation of the stride of the convolutional kernel or with pooling [34]. Convolutional layers have benefits over fully-connected layers when using images as inputs, since they apply the same kernel across the entire image, reducing the number of parameters required while also making the network more robust to the translation of features in the image [34].

In this section, the machine learning model loss function and architecture will be explained, as well as the image pairing required for this method to work. The interpretation of the model outputs is explained, and how they are translated into the process maps is presented later in the results section.

3.1. Invariant information clustering

The chosen ML algorithm uses the Invariant Information Clustering (IIC) loss function detailed in the work of Ji et al. [35]. This approach will be summarised here to provide context, but the full mathematical explanation and benchmarking can be found in the original work of Ji et al. [35]. In short, the IIC algorithm seeks to group images based on their similarity without any prior knowledge of labels or categories. This is achieved by training a neural network (NN) to maximize the mutual information between the output probability distributions of a positive pair of images (different images that should be assigned the same cluster assignment). Other methods may use negative pairs (images that should be assigned different categories) or even triplets (original images, positive version and negative version) as the input and compare them differently. The details of how the image pairing is performed in this work are given in Section 3.4. In the original work of Ji et al., [35], synthetic image augmentation is used to create the image pair to avoid the lengthy process of manually pairing different image samples.

The neural network is trained to learn a mapping $\Phi : \mathcal{X} \rightarrow \mathcal{Y}$, where $\mathcal{X}$ is a joint probability distribution containing both the paired data sample, $(\mathbf{x}, \mathbf{x}')$, and $\mathcal{Y}$ is the set of possible cluster assignments, $\mathcal{Y} = \{1, \ldots, C\}$, being C the total number of clusters. As is implied in the name of this method, the goal is to identify what is common, or *invariant*, between the images and discard details that do not contribute to the final classification. This is naturally achieved using a finite number of cluster assignments as a bottleneck in the network. The loss function used in the IIC algorithm [35] is the maximisation of mutual information between the pair of samples:

$$\max_{\Phi} I(\Phi(\mathbf{x}), \Phi(\mathbf{x}')), \tag{1}$$

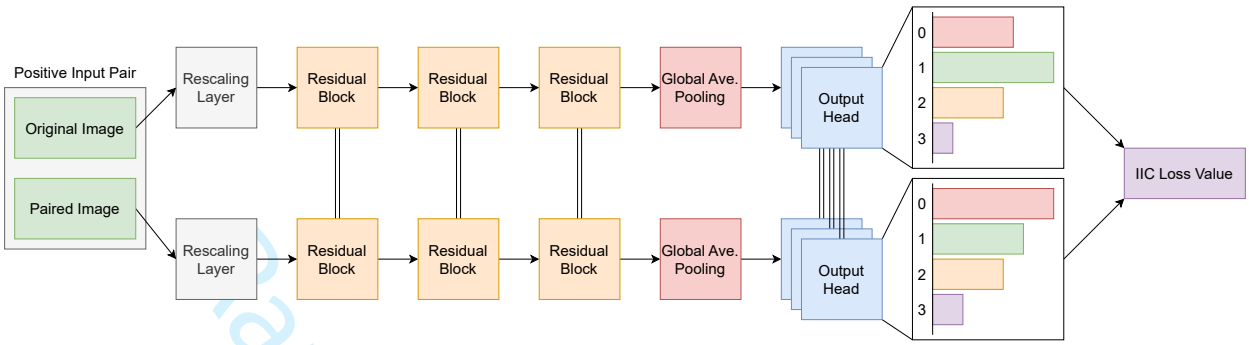

Figure 5: Diagrammatic representation of the neural network used in this work. The double lines indicate weight sharing between the corresponding layers. *Authors own work.*

where $I(\Phi(\mathbf{x}), \Phi(\mathbf{x}'))$ is the mutual information between encoded variables. Maximising the mutual information is essentially maximising the predictability of one sample based on the other. The neural network is terminated with a SoftMax activation layer, that provides a discrete probability distribution across the chosen number of clusters, $\Phi(\mathbf{x}) \in [0,1]^C$. More formally, this can be written as $P(z = c|\mathbf{x}) = \Phi_c(\mathbf{x})$, where z is a cluster assignment for input $\mathbf{x}$ over C-classes. The conditional joint probability distribution for 2 inputs $\mathbf{x}$ and $\mathbf{x}'$ given cluster assignments z and z' is given as $P(z = c, z' = c'|\mathbf{x}, \mathbf{x}') = \Phi_c(\mathbf{x}) \cdot \Phi_{c'}(\mathbf{x}')$. With marginalisation across a batch of inputs during training, the joint probability distribution [35] is represented by a $C \times C$ matrix $\mathbf{P}$:

$$\mathbf{P} = \frac{1}{n} \sum_{i=1}^{n} \Phi(x_i) \cdot \Phi(x_i')^{\mathrm{T}} \tag{2}$$

$\mathbf{P}$ is symmetrised by using $(\mathbf{P}+\mathbf{P}^{\mathrm{T}})/2$, and then substituted into the formula for mutual information [36] to calculate the mutual information in the following expression [35]:

$$I(z, z') = I(\mathbf{P}) = \sum_{c=1}^{C} \sum_{c'=1}^{C} \mathbf{P}_{cc'} \cdot \ln \frac{\mathbf{P}_{cc'}}{\mathbf{P}_c \cdot \mathbf{P}_{c'}}, \tag{3}$$

$\mathbf{P}_{cc'}$ represents $P(z = c, z' = c')$, and $\mathbf{P}_c$ and $\mathbf{P}_{c'}$ are the marginals, calculated by summing across the rows and columns of $\mathbf{P}$, respectively.

Now that the loss function I has been defined in Equation (3), the neural network can be trained across epochs using a standard optimisation algorithm, such as stochastic gradient descent. This will be explained in the following section.

3.2. Machine learning model architecture

The neural network used in this work was created from scratch using the TensorFlow Python API [37]. Figure 5 provides an overview of the network architecture: a rescaling layer, 3 residual blocks [38], and a global average pooling layer [39]. The rescaling layer maps the pixel values in the images to between 0 and 1 by dividing by 255. The 3 residual blocks perform the convolution operations, transforming the images using trainable convolution kernels and extracting the key features. Each residual block contains 2 3x3 convolution layers with 16 filters each, and the image resolution is halved with each block by striding the convolutional filters by a factor of 2. The skip connection in the residual block uses a 1x1 convolution with stride to match the image dimensions. More information for these residual blocks is given in Appendix A.1. Finally, a global average pooling layer transforms the image feature arrays into a single vector to be interpreted by the fully connected prediction heads.

Once the image has passed through the convolution part of the network, it comes to the prediction heads, where the class of the images will be assigned. Each output head contains 2 hidden dense layers of 32 neurons, each with a batch normalisation layer preceding the ReLU activation function. These heads contain a final dense layer whose length is equal to the number of clusters desired, in this case, 4. This final layer terminates with a SoftMax activation

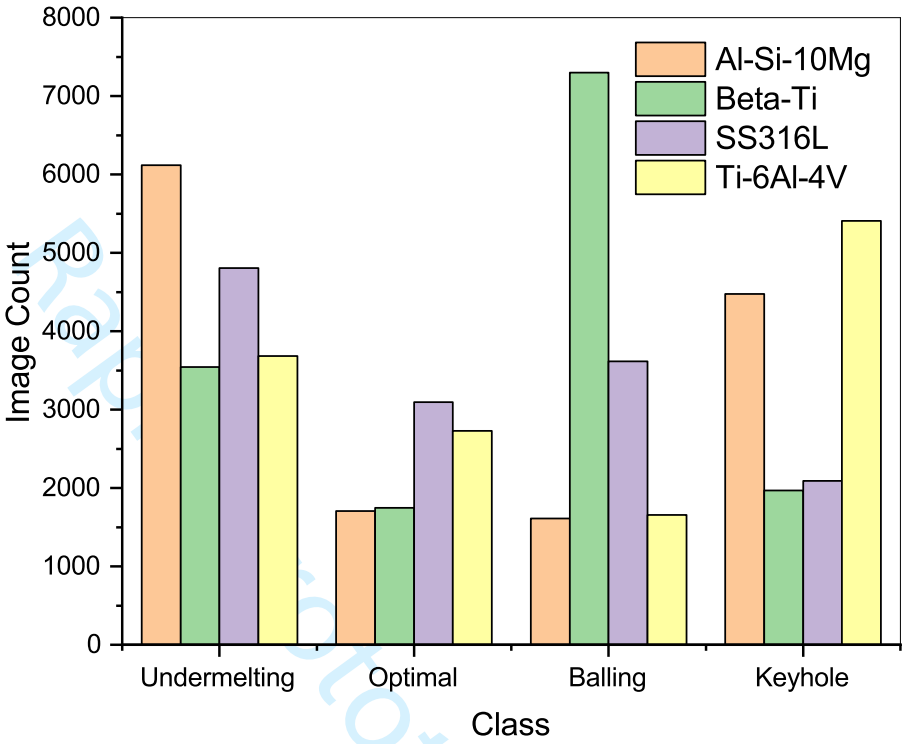

Figure 6: The distribution of classes in the 4 alloys. *Authors own work.*

function, providing a discrete probability distribution across the 4 categories. More information about the architecture is given in Appendix A.2.

In the original work of Ji et al. [35], the authors connected multiple output heads to the CNN, each with the same number of output clusters. The overall loss is then calculated as an average of the losses from the different output heads, and the head with the lowest loss after training is selected as the final prediction head. In this work, we take the same approach and connect 3 output heads to the CNN, each with an output size of 4 clusters. The models were trained for 20 epochs, using the Adam [40] optimiser from Tensorflow [37] with an initial learning rate of 1e-3 with exponential decay and a batch size of 32. All training was performed on local hardware with GPU acceleration.

3.3. Dataset balancing

In Section 2.2, we discussed how the ground truth morphology of each track was determined. Using this ground truth information, we extract the distribution of classes in the alloy datasets. Figure 6 shows the distribution of the classes in the dataset. There is a large variation in the number of images per class in each alloy. During initial testing of the model, we found that the classes with larger numbers of images tended to have a higher prediction accuracy, while those classes with smaller numbers of samples were often ignored by the model. This was particularly problematic for the Al-Si-10Mg alloy, which has many undermelting and keyhole tracks relative to the number of optimal and balling samples. This resulted in the model failing to predict that any tracks were optimal or balling and instead, it assigned them to either undermelting or keyhole.

To combat this, we sub-sampled these datasets to even the number of samples in each class. For each alloy, the class with the lowest number of samples acted as the target for the sub-sampling. Then, for the remaining 3 classes, we randomly subtracted the difference, leaving each class with the same number of samples. To ensure even representation for each track, this sampling was performed on a track-by-track basis, with an equal number of samples being randomly removed from each. The samples that were removed from the databases were used as the testing data for each alloy. A further 20% of images were removed to ensure that all classes were represented in the testing dataset.

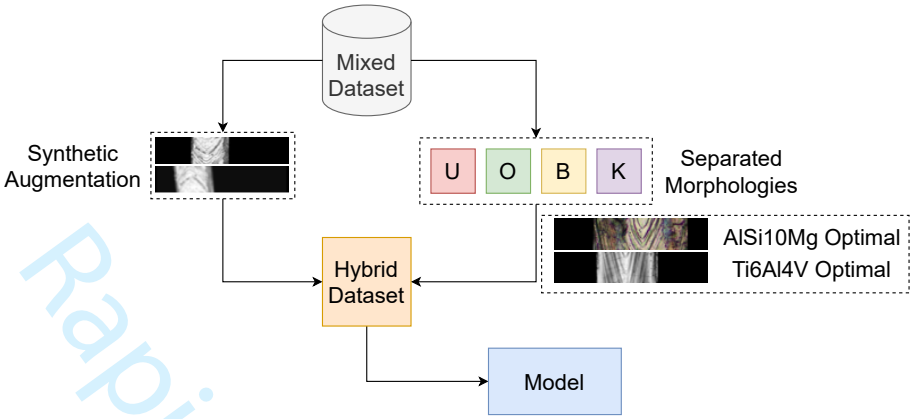

Figure 7: The construction of the hybrid dataset used to drive alloy agnostic behaviour. *Authors own work.*

3.4. Hybrid training strategy

In the original work of Ji et al.[35], the authors used synthetic image augmentation to modify the original images to generate a positive pair. The image transformations used were selected to sufficiently modify the image to create a separate sample while maintaining the original structure of the data. Such transformations included rotation, translation, contrast and brightness changes. This was done since the dataset was immense, and manual labelling would have been a time and labour-intensive task. For our dataset, we know the ground truth data for the images from the cross-section analysis in Section 2.2, so pairing images based on their morphology is a relatively trivial task. However, given the success of using image augmentation in the original implementation, we wanted to test the performance of our models when using fully synthetic augmentation, fully labelled pairing, and then a combination of these two. We define a ratio r to quantify the level of manual pairing in our dataset, where $r = 0$ means the data is unlabelled and uses augmentation, and $r = 1$ means that all image pairs are constructed using the ground truth data.

As discussed in the introduction and later in Section 2.2, there exists some uncertainty in the ground truth data, since the cross-sectional images taken may not represent the melting mode of the entire track. Using fully supervised training would assume the ground truth data is perfect, so by having a mixture of supervised and unsupervised training, this uncertainty can be instead used as a guide in the training process, rather than the absolute truth.

To manually pair our images, we first need to organise the samples into 4 subsets according to their morphology type, but independently of their alloy system. During training, when an image is sampled from a particular morphology type, its pair is also randomly sampled from the same morphology type. This random sampling can result in pairs of the same alloy, or different alloys, but both with the same track morphology. To create the mixed dataset where $0 < r < 1$, a fraction r images are taken from the original unlabelled set, and organised into morphology types as described above. During training, images are sampled from both the labelled and unlabelled datasets, resulting in batches containing both random augmentation and manual pairing, with a likelihood that depends on r. With this hybrid dataset approach, the model learns to understand the unseen data resulting from the image augmentation, and also learns to ignore the alloy difference, focusing instead on what features each morphology class has. This approach is summarised in Figure 7. In the results, we will present the effect of varying r on the accuracy of the process maps.

3.5. Matching clusters to ground truth morphologies

The model proposed in this work terminates in a SoftMax activation layer, providing the probability distributions for each image sample across the given number of clusters. The final cluster assignment for an image sample can be extracted from this probability vector using an ArgMax function, which returns the cluster with the highest probability. This probability represents the model's certainty of the cluster assignment out of all the possible clusters in a particular prediction head. As discussed in Section 2.3, the image slices fed into the network are small portions of each larger track. To determine the final melting mode of the entire track, the cluster assignments for each of the slices forming that track are counted, and the cluster with the largest number of assignments is selected as the overall assignment of the track. This method allows for a few misclassifications without significantly affecting the results.

	Case 1	Case 2	Case 3	Case 4
Training	β-Ti SS316L Ti-6Al-4V	Al-Si-10Mg SS316L Ti-6Al-4V	Al-Si-10Mg β-Ti Ti-6Al-4V	Al-Si-10Mg β-Ti SS316L
Validation	Al-Si-10Mg	β-Ti	SS316L	Ti-6Al-4V

Table 2: 4-fold validation training and testing configurations. *Authors own work.*

The unsupervised model does not directly assign names to the different clusters of tracks, it just groups the most similar images. To associate these clusters to each melting mode seen in the ground truth data, the Jaccard Index (J) is calculated between the ground truth data extracted from the cross-section images (A) introduced in the previous subsection, and the predicted labels (B):

$$J(A, B) = \frac{A \cap B}{A \cup B} \tag{4}$$

The Jaccard Index varies from 0 to 1, with 0 indicating that the intersection between the 2 sets is empty, and 1 representing a perfect match between an experimental cluster and a predicted cluster. The highest Jaccard score is used to assign morphologies to the clustered groups.

3.6. 4-fold validation process

in this work, we want to create a model that is robust and alloy-agnostic. We will demonstrate this using a 4-fold validation process, where each alloy from our training dataset is removed to be used as a validation alloy, unseen by the model during training. This is summarised in Table 2.

Taking this approach allows us to explicitly test the generalisability of the model depending on the training conditions. Since we are changing the training alloys each time, we expect the accuracy of the predicted process maps to change also. In Section 3.4, we discuss how the r ratio of the data can be changed to modify how much of the image data is paired using ground truth data. In this work, we use 3 values of r, 0, 0.5, and 1.0. For each value, the 4-fold validation training process is used to assess how the model accuracy changes as a function of r. These results are discussed in Section 4.

As discussed in Section 3.3, the training datasets were balanced before training to ensure an even representation for each morphology class. The removed images are used as the testing data to validate the accuracy of the training alloys.

4. Results and discussion

In this section, we will begin by discussing the effect of changing the r ratio on the model accuracy and generalisability, before talking about the results of the 4-fold validation study in more detail. This 4-fold validation process helps us to understand how robust the model is to different compositions of training data, and characterise the generalisability of the model under these changing conditions.

4.1. Effect of synthetic to manual pairing ratio (r)

As described in Section 3.6, we used the 4-fold training process to investigate how the model behaviour changes with different alloy combinations in the training data, but also with differing r ratios. The effect of r on the average final model accuracy and the generalisation accuracy is shown in Figure 8.

Figure 8 demonstrates several important behaviours of the model. Primarily, the training accuracy is higher than the generalisation accuracy in all cases, which is to be expected. Second, the training and generalisation accuracy are both highest when the ratio $r = 0.5$, or the training pairs are generated by synthetic image augmentation and manual pairing. And finally, the spread in the data is also lowest when $r = 0.5$. We can therefore conclude that using a model trained with 50% synthetic pairs and 50% manual pairs is the most robust (with the lowest spread in the accuracy) and generalisable (with the highest generalisation accuracy). The data from the graph is included in the appendix,

For the remainder of the results in this paper, we will focus on the results for when $r = 0.5$.

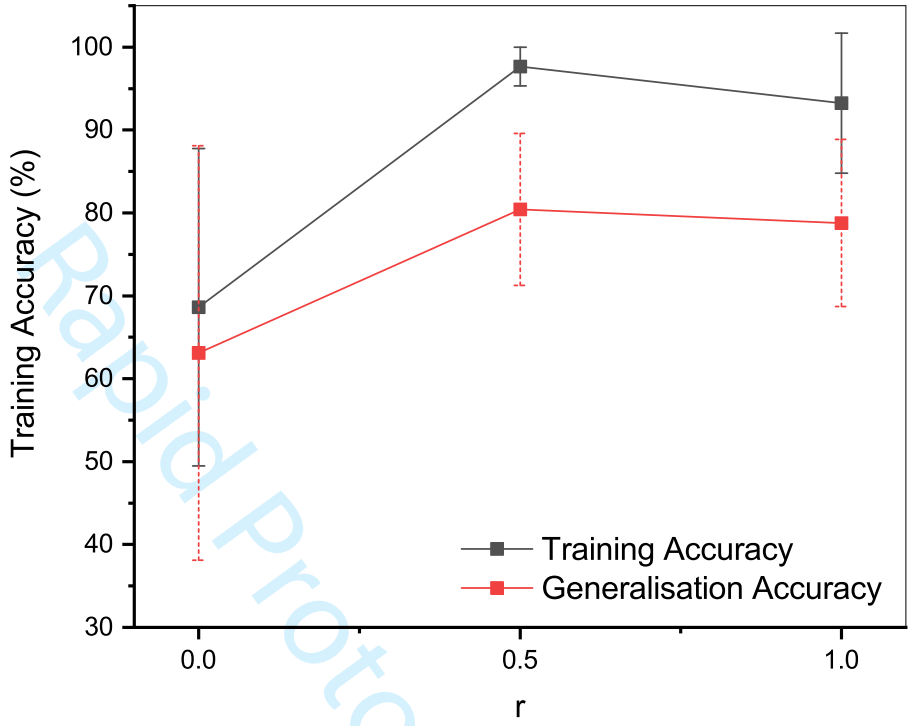

Figure 8: Average training and generalisation accuracy vs ratio r, the fraction of labelled data. The error bars show the range of accuracy for each value of r. *Authors own work.*

4.2. Clustering performance

In Section 4.1, we discussed how the model is more robust and generalisable when the training data contains a mixture of synthetically generated image pairs and manually paired samples. In this section, we will present the results of the 4-fold analysis in more detail, discussing both the accuracy of the model prediction on the training alloys, as well as the unseen alloys for each case.

The complete accuracy results for the 4-fold analysis of the model with each case and each of the 3 prediction heads are shown in Table 3, along with the average accuracy. What we would expect to see for a robust model would be for a certain case, the 3 prediction heads result in the same accuracy, in some percentage points. Also, we would expect that the alloy not contained in the training dataset (referred to as generalisation alloys from here on) has a lower prediction accuracy compared to the 3 training alloys.

The predictions in Table 3 reflect these expectations. For each model, the process maps for the training alloys are predicted consistently with an accuracy equal to or greater than 90%, and this accuracy is consistent across each prediction head per model. For the generalisation alloys (indicated with a * in Table 3), the prediction accuracy is lower, which is to be expected. Table 3 also shows that the accuracy of prediction on the generalisation alloys does vary depending on the training data.

4.2.1. Case 0 - Al-Si-10Mg validation

For case 0, the Al-Si-10Mg is the generalisation alloy. This case demonstrates the model's performance on a new alloy that behaves differently from the training alloys since the conductivity of the Al-Si-10Mg alloy is higher than in the titanium alloys and the stainless steel. It is also well known that aluminium alloys, including Al-Si-10Mg, are considered to have high reflectivity [22, 41], and this adds challenges to their manufacture with LPBF. Patel et al.[42] in their work discuss how the high conductivity and reflectivity properties of Al-Si-10Mg cause a rapid onset of keyhole mode melting in LPBF. This makes finding a stable processing window challenging for this material. This can be seen clearly in Figure 9, in the true map. The optimal region only contains 5 out of the 40 total tracks.

Training Case	Prediction Head	Accuracy (%)				
		Al-Si-10Mg	β-Ti	SS316L	Ti-6Al-4V	Average
0	0	75.0*	96.0	92.5	95.0	89.6
	1	72.5*	96.0	90.0	97.5	89.0
	2	77.5*	96.0	90.0	97.5	90.3
1	0	97.5	96.0*	100.0	97.5	97.8
	1	97.5	92.0*	100.0	97.5	96.8
	2	97.5	92.0*	100.0	97.5	96.8
2	0	100.0	100.0	82.5*	97.5	95.0
	1	100.0	100.0	82.5*	97.5	95.0
	2	100.0	100.0	82.5*	97.5	95.0
3	0	100.0	100.0	97.5	75.0*	93.1
	1	100.0	100.0	95.0	75.0*	92.5
	2	100.0	100.0	95.0	62.5*	89.4

Table 3: Prediction accuracy for each output head of the model when trained using a specific case of training alloys (from Section 3.6). * indicates the generalisation alloy in that case. $r = 0.5$ for these results. *Authors own work.*

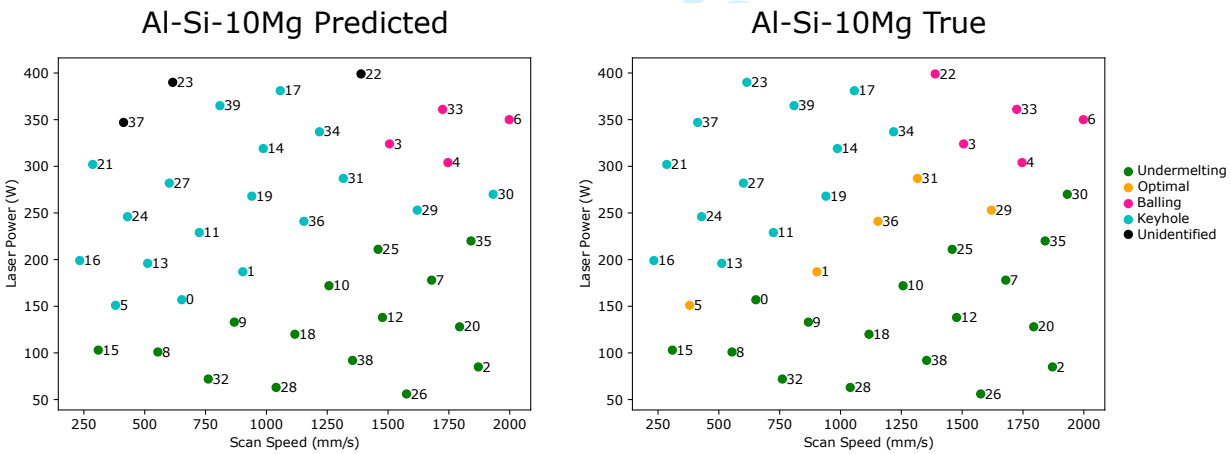

Figure 9: Predicted process map for the Al-Si-10Mg for Case 0, from prediction head 0 and the true map for the Al-Si-10Mg. The numbers are a reference to the printing parameter set. *Authors own work.*

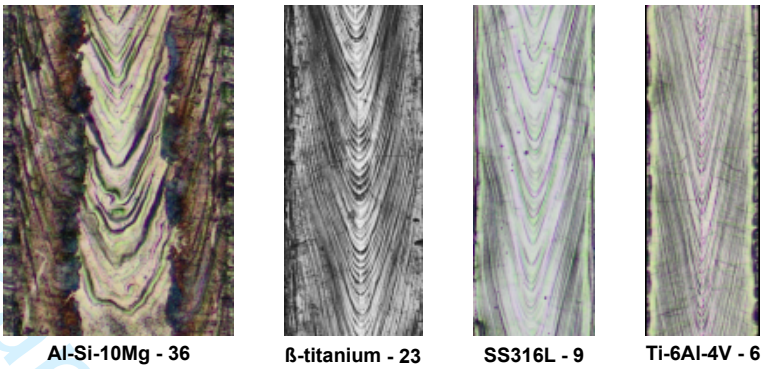

Figure 10: Comparison between optimal tracks for the 4 tested alloy systems. The scale is constant for all of the tracks. *Authors own work.*

Figure 9 shows the predicted process map for the Al-Si-10Mg for training case 0, using prediction head 0, next to the true Al-Si-10Mg process map. We can immediately see that the optimal group is not predicted. The likely reason for this is that the information learned from the provided training alloys does not transfer well to the Al-Si-10Mg. Instead, the optimal tracks have been classified as keyhole. This suggests that for Al-Si-10Mg, the optimal tracks are most similar to the keyhole tracks. Given that the conductivity of Al-Si-10Mg is significantly different to the training alloys, it is expected that the prediction accuracy suffers. By including an alloy (or alloys) that is more similar to the Al-Si-10Mg in the training dataset, the generalisation accuracy should increase.

Despite the lack of the optimal region in the predicted processing map, the undermelting and balling groups are predicted well, suggesting there is enough information in the training alloys to predict these morphologies in the Al-Si-10Mg. This suggests that the optimal tracks of the Al-Si-10Mg exhibit features not observed in the optimal tracks of the alloys in the training data. Figure 10 contains images showing optimal track sections for each of the tested alloys systems. Just by visual comparison, it is clear that the Al-Si-10Mg track is very different to the other 3. There is a lighter central region, flanked by darker regions. The central region also has thicker and more irregular fish-scale like features along its length. This demonstrates just how different the melting behaviour is for the Al-Si-10Mg alloy in comparison with the others, and it is therefore understandable why the model does not predict the appearance of the optimal melting zone.

This result is promising because we show that the model gains sufficient information from 3 alloys to predict the behaviour of a very different alloy system with reasonable accuracy. As is usually the case with data-driven approaches, we predict that with a wider range of alloy systems and more data samples, the model will generalise to new alloys better, and we would observe an improvement in this case.

4.2.2. Case 1 - β-titanium validation

Case 1 represents the prediction of an alloy similar to those contained in the training data. The β-titanium alloy is a titanium alloy, so it will behave similar way to the Ti-6Al-4V, which is part of the training dataset.

After training, case 1 has the highest generalisation accuracy and overall prediction accuracy of the 4 different cases. The high generalisation accuracy suggests that the model has learned sufficient information from the Al-Si-10Mg, SS316L and Ti-6Al-4V alloys to predict the behaviour of the unseen β-titanium with high accuracy.

Figure 11 shows that the model has learned all the necessary features to make an almost perfect prediction of the β-titanium alloy process map without having seen it during the training process. Only 2 of the tracks have been misclassified, meaning the optimal zone in the predicted map is smaller than in the ground truth data. This shows that the model has made a conservative prediction in the process map for this alloy system. Importantly, the presence of an alloy that is not similar (the Al-Si-10Mg) does not appear to negatively impact the generalisation performance.

This process map shows that the 3 alloys contained in the training dataset contain sufficient information for the model to infer the behaviour of the new β-titanium alloy, and predict the process map with a high level of accuracy. We expect that with more training data, the model would predict the process map for the β-titanium alloy perfectly.

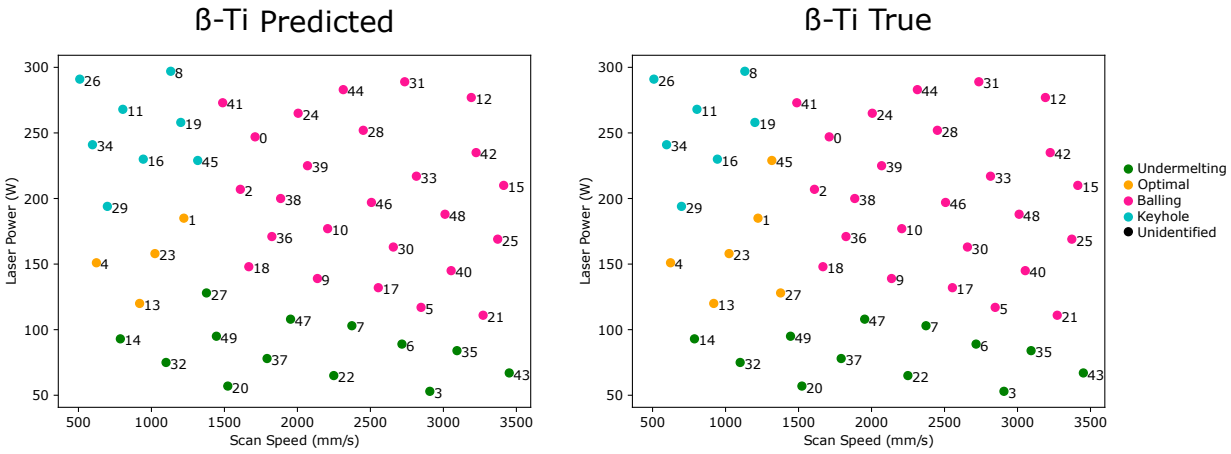

Figure 11: Predicted process map for the β-titanium for Case 1, from prediction head 0 and the true map for the β-titanium. The numbers are a reference to the printing parameter set. *Authors own work.*

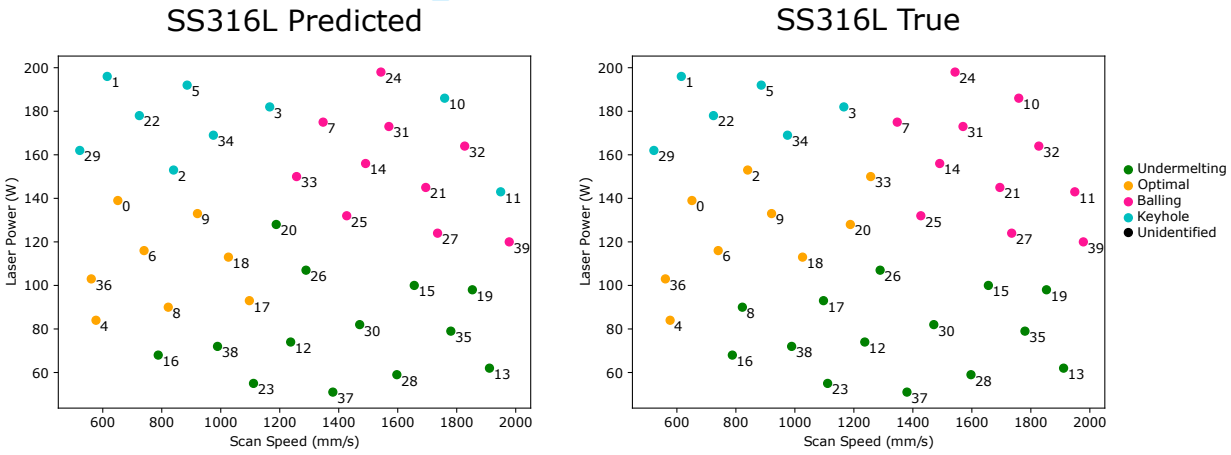

Figure 12: Predicted process map for the SS316L for Case 2, from prediction head 0 and the true map for the SS316L. The numbers are a reference to the printing parameter set. *Authors own work.*

4.2.3. Case 2 - SS316L validation

Case 2 represents another situation where the generalisation alloy is a different alloy system to those contained in the training data, the SS316L stainless steel. However, this time the thermal conductivity is similar to the Ti-6Al-4V and the β-titanium alloys. We therefore expect the melt behaviour to be more accurately predicted by the trained model. For case 2, all 3 output heads predict the processing maps with the same 82.5% accuracy for the generalisation alloy. The predicted process map from head 0 is shown in Figure 12.

Figure 12 shows that the model has learned the underlying features of the images well enough to provide a good representation of the process map. There are some discrepancies, however, including some confusion between the optimal and undermelting regions, as well as 2 balling tracks being predicted as keyhole, but overall the performance is promising.

4.2.4. Case 3 - Ti-6Al-4V validation

The final case, Case 3 also shows lower generalisation accuracy, but this time for Ti-6Al-4V as the generalisation alloy. Figure 13 shows the process map for head 0 in this case next to the ground truth Ti-6Al-4V process map.

In this figure, we can see that the 4 morphology groups are present, with no tracks being unidentified. This shows

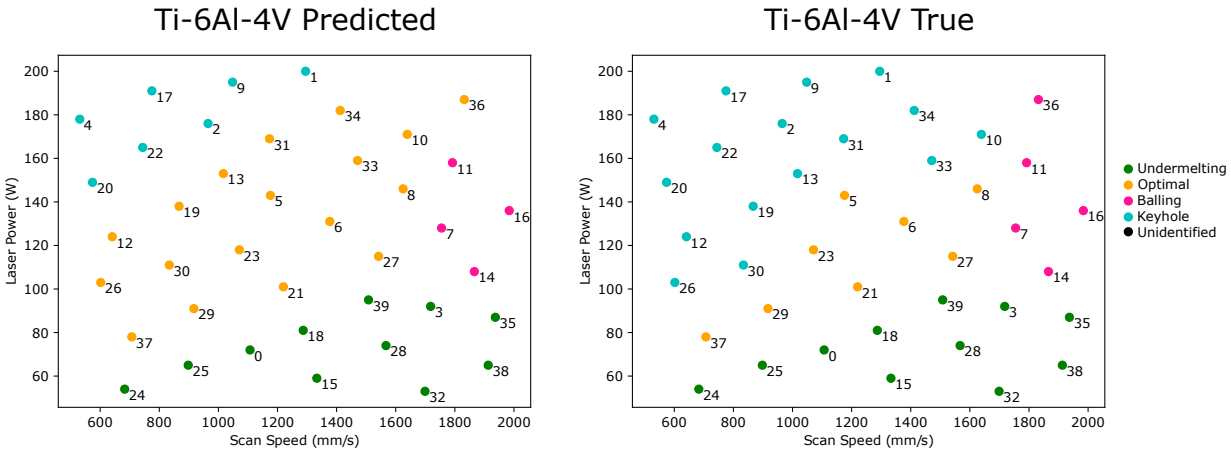

Figure 13: Predicted process map for the Ti-6Al-4V for Case 3, from prediction head 0 and the true map for the Ti-6Al-4V. The numbers are a reference to the printing parameter set. *Authors own work.*

that the model can learn some information from the training alloys and make a reasonable prediction for the behaviour of the Ti-6Al-4V. This map's primary source of error comes from the interface between the optimal and keyhole tracks. A significant number of the keyhole tracks have been misclassified as optimal, which enlarges the predicted optimal processing region for this alloy. If used in practice, this would mean that some printed samples will likely exhibit keyhole defects. This behaviour is likely due to the melting morphology of the tracks from the upper surface view exhibiting a gradual change from optimal to keyhole. The model has not been able to exactly pinpoint the transition from optimal to keyhole for this alloy system, and instead incorrectly classifies some keyhole mode tracks as optimal.

Compared to the model's performance under case 0, the generalisation performance for the Ti-6Al-4V is lower than for the Al-Si-10Mg when the accuracy metric shown here is used. However, this should not be looked at in isolation. Instead, it should be noted that all 4 morphology groups are present in the Ti-6Al-4V process map, whereas they are not in the case of the Al-Si-10Mg. For this reason, the process map for the Ti-6Al-4V is much more useful and of higher quality than that for the Al-Si-10Mg.

We expected the accuracy of this process map to be similar to that of the β-titanium and 316L stainless steel maps since the model was able to learn useful information from the Ti-6Al-4V and apply it to the generalisation alloys in those cases. This learning does not appear to flow both ways, meaning that the information gained from the β-titanium and the 316L stainless steel is less useful in predicting the process map for the Ti-6Al-4V.

4.3. Consideration of powder properties and other scanning variables

As described in Section 2.1 the single tracks printed in this work do not use a powder layer and are the weld tracks left behind after the laser scans the bare surface of the sample. The reason for this is that the objective of this work is to create a classifier of melting modes based on the top-down view of the scanned tracks to enable process engineers to determine the printability of the alloy system before powder production. Hence, the decision was made to remove the powder as a variable in the experiment and focus on developing an experimental and computational pipeline for image recognition across different alloy systems. Removing the powder layer requirement allows process engineers to use solid ingots to determine the processing window, which may vary in size and shape. All that is required in this case is a flat scanning surface, which has been ground from smooth. Once the printability of the proposed alloy system has been verified using the desired hardware, process engineers can continue with parameter optimisation using the powder feedstock.

It has been widely shown in the literature that the powder properties, such as layer thickness and particle size distribution, do play a significant role in determining the final density and mechanical properties of the solidified material, as highlighted experimentally in the work by Gor et al. [43] and computationally in the work by Cao et al. [44]. In fact, after scanning speed and laser power, the work by Gor et al. [43] states that layer thickness is the next most influential parameter in the density of the final printed material. Ziri et al. [45] performed an in depth study on the

relation between the input energy density, powder properties and resulting part properties. They determined that the particle size distribution of the feedstock powder will affect the types of defects that form at constant energy densities and can even affect where keyhole and lack of fusion boundaries occur, narrowing the optimal processing window.

The method proposed in this work is applicable to images of single tracks printed with a powder feedstock layer as well, provided the training images are available. The main limiting factor in this case would be the availability of a wide variety of different powder materials to generate diverse training data, and ensure the reliable deposition of the powder layer of a specific thickness. One also cannot neglect the work involved with changing the feedstock material between experiments, assuming the same machine is to be used for all experiments. Future work would need to address some of these challenges in obtaining sufficient experimental data in a time-efficient manner, but the proposed method would be suitable.

The proposed method also neglects other scanning variables, such as hatch spacing. While the model does not directly output this, the top-down images can be measured to determine their width w, and then the hatch spacing s, can be estimated using an overlap fraction o, shown in Equation (5). A Gaussian process model could be constructed from the printed tracks to estimate the track width at an arbitrary power and scan speed. This approach is demonstrated in Section 4.5, where o is set to 0.3.

$$s = \frac{w}{2} * (1 - o) \tag{5}$$

The optimal layer thickness for printing cannot be estimated using the current proposed framework, as the process relies on top-down images. This is the compromise for considering more of the track length in the prediction of the melting mode. However, one possibility would be to construct a surrogate model of the track depths from a design of experiments. This would require measurement of the depth for each of the tracks, so a cross-section view or a CT scan. Then, this model can be queried for an arbitrary set of power and scanning speed to determine the predicted track depth, and the appropriate layer height could be estimated like the hatch spacing using Equation (5). A similar approach is discussed in Section 4.5.

4.4. Improving the generalisability of the model

One of the key objectives within this work is to create a model that can generalise well to new alloy systems, given that the purpose is to enable process engineers to quickly determine a process map for a novel alloy. This work has already presented numerous strategies to extend the capabilities of the model outside of the training data. Primarily, this work introduced a novel hybrid training strategy in Section 3.4, which presented a way to combine both supervised and unsupervised learning that outperforms each in terms of training and generalisation accuracy, as well as with reduced error, as discussed in Section 4.1. Besides this, a number of commonly seen design tactics in neural network architecture were used to improve the stability of training and introduce some regularisation, including batch normalisation and residual blocks. The complexity of the model was also kept lower to minimise the risk of overfitting. A 4-fold validation process was used to quantify the model's performance across a range of training cases, to observe how well the model could generalise to different alloy systems.

Despite these efforts, the results for the Al-Si-10Mg specifically showed that when the model is trained using alloy systems that differ greatly from the alloy system used for inference, performance is poor. In this case, the optimal zone for the Al-Si-10Mg was not detected at all. In Section 4.2.1 we hypothesis that there is not enough information contained in the three training alloys to predict the optimal melting zone for the Al-Si-10Mg alloy, as the high reflectivity causes the melting behaviour of this alloy system to significantly differ from the training data. We observed that the other three melting modes were predicted well, suggesting these are more material-independent. Luckily, there already exist a number of methods to improve neural network performance within the machine learning community, and in future work, some or all of these could be implemented to improve the model's performance in predicting this optimal melting mode for the Al-Si-10Mg. The authors believe that the method which may yield the best results in this case is to perform feature extraction on the input images prior to their use as input for model training. By distilling the essence of each image into a smaller set of features, the model is able to focus its training and make better utilisation of the available network capacity. This would need a study to determine which features are the most important to extract.

Another route to improving the model performance would be to perform an optimisation study on the architecture itself. This would determine certain architectural parameters, for example the number of convolutional filters in the

Sample	True Morphology	Predicted Morphology	Power (W)	Scanning Speed (mm/s)	Hatch Spacing (μm)	Density (%)
A	Optimal	Optimal	110	1300	68.00	99.96
B	Optimal	Optimal	80	800	68.00	99.96
C	Optimal - Keyhole	Optimal	160	1600	65.00	99.95
D	Keyhole	Optimal - Keyhole	160	950	95.00	99.95
E	Keyhole	Optimal - Keyhole	180	1200	78.00	99.97
F	Keyhole	Optimal	110	600	85.00	99.81
G	Balling	Balling	120	1800	55.00	99.96
H	Balling	Balling	140	1900	58.00	99.98
I	Balling	Optimal - Balling	180	1950	60.00	99.99
J	Undermelting	Undermelting	70	1400	45.00	96.01
K	Undermelting	Undermelting	60	800	49.00	98.99
L	Undermelting	Undermelting	80	1900	40.00	95.77

Table 4: The printing parameters used to fabricate the 3D test samples, with the true and predicted melting mode, and the optical density. *Authors own work.*

residual blocks, or the number of layers in the output heads. By concentrating model parameters where they are most useful, the complexity of the network can be reduced, and this may also result in faster training, and less likelihood of overfitting.

The final method for improving the generalisability of the model is to simply use more training data, and more alloy systems in that dataset. In this work, we have seen that where a similar alloy is present in both the training and testing data, the generalisation is good. So for predicting the Al-Si-10Mg process map, if an alloy system with a similar melting behaviour were present, we would expect the model to perform better, and identify the optimal melting zone.

4.5. Application of the framework in 3D

For the completeness of this study, a set of 3D samples was fabricated to determine whether the observed powder-free, single-track morphologies were also seen in multi-layer manufacturing, and thus determine the potential of such a neural network in the prediction of the final component's printability. These samples were cubes with a 5mm edge length, using Ti-6Al-4V powder feedstock. For each zone on the true processing map, 3 parameter sets were selected, giving 12 processing conditions. The layer thickness was fixed at 30μm, and a hatch-only strategy was chosen (no contour scan). A Gaussian Process surrogate model tuned using the measurements extracted from the track cross-sections was used to predict the track widths for these new parameter sets and as a result, the hatch spacing. The hatch spacing for each cube is calculated to ensure a 30% overlap between adjacent scan tracks. The printing parameters are summarised in Table 4 below, where the measured optical density is also given, the true morphology from the cross-sectional analysis, and the predicted morphology from the neural network.

Cross-section images for each parameter set can be seen in Figure 14. Table 4 shows that the optical density for cubes printed outside of the undermelting region exceeds 99.8%. Looking at Figure 14, the undermelting cubes have a large number of irregularly shaped pores which are aligned perpendicular to the printing direction (vertical), which is typical for lack of fusion porosity [9, 29, 46]. The optimal and keyhole cross-sections observed in Figure 14 contain a combination of irregular-shaped pores as well as smaller, spherical or elliptical pores, which occur when gas is trapped in the alloy during solidification [9, 29, 46]. The cubes printed with keyhole parameters exhibit higher porosity when compared to the optimal cubes.

Studies conclude that when scanning with high scanning velocities (balling mode), the unstable melt pool results in a poor surface finish, which prevents quality powder spreading and can cause irregularly shaped defects [29]. In this work, our samples printed with the balling morphology tracks did not exhibit a significantly lower optical density when compared to the optimal baseline, suggesting that for our parameters and equipment, the balling effect seen in the single tracks was not severe enough to disrupt the printing process.

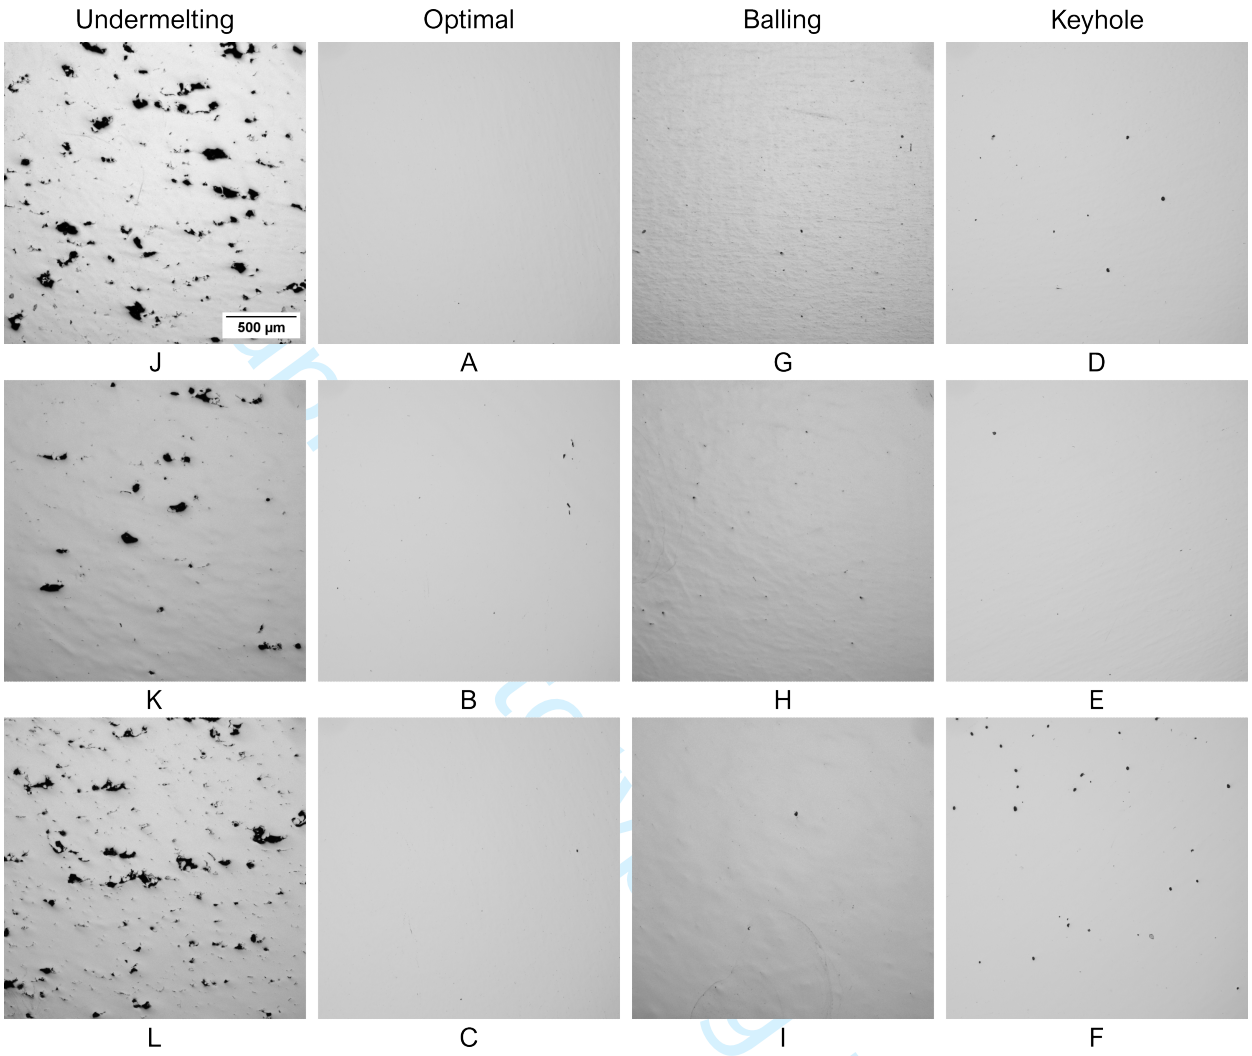

Figure 14: Cross-section images from the cube samples printed in the different melting morphologies as predicted by the ML model; optimal (A, B, C), keyhole (D, E, F), balling (G, H, I) and undermelting (J, K, L). The printing direction is from bottom to top for all images. The images are taken from the samples printed with the parameters in Table 4. All images have the same scale. *Authors own work.*

This analysis suggests that the powder-free single-tracks are a useful indicator for the final part porosity and therefore quality. This means that having a neural network capable of correctly identifying the morphology type of powder-free single tracks is a valuable tool for determining the final printing parameters for a component. In future work, it would be valuable to expand the parameter set to trigger the balling and keyhole effects to show more clearly in the printed samples. The current parameter range (outside of the undermelting tracks) provides cube samples with high optical density, and we do not observe the keyhole or balling effects.

The process provided in this section suggests how the proposed machine learning framework could be used to determine a suitable processing window for a given alloy. This will be briefly summarised:

1. Scan several single tracks on a bare alloy substrate with varying laser powers and scanning speeds
2. Use the proposed trained model to identify the desired melting mode for further processing
3. Calculate the hatch spacing for a parameter set by measuring the printed track widths and ensuring a suitable overlap (0.3 is used in this work) using Equation (5).
4. Once powder material is available, fabricate test coupons using the determined power, scan speed, and hatch spacing parameters. The layer height may be chosen as a standard number (30µm in this study) or result from

further parameter optimisation work. Alternatively, measure the track depths and apply an overlap to calculate an appropriate layer height using Equation (5).

5. Conclusions

In this work, we used a convolutional neural network to automatically generate the process maps for 4 different alloy systems, Al-Si-10Mg, a β-stable titanium alloy, SS316L stainless steel, and Ti-6Al-4V. For all cases tested, the process maps for the training materials are predicted with accuracies exceeding 90%. We demonstrated the effect of changing training conditions on the generalisation accuracy of the model, and Our results show that the generalisation performance is generally better for materials that are similar to those contained in the training set. When the training set contains less variation in alloy systems, the generalisation accuracy tends to be lower. If the material is too different to the training data, the model may not be able to detect the presence of all 4 morphology types, as was seen in the case of the Al-Si-10Mg.

The best generalisation performance was seen when the β-titanium alloy process map was predicted using the Ti-6Al-4V, 316L stainless steel and the Al-Si-10Mg as training data. The accuracy of the process map ranged between 92.0 and 96.0%. The worst performing cases were for the Al-Si-10Mg, where the optimal processing zone was missed, and the Ti-6Al-4V, where the boundary between the optimal and keyhole modes was incorrectly predicted. The accuracies were 72.5-77.5% and 62.5-75.0% respectively for these alloys.

The proposed method considers the entire length of the track when determining the melting mode, instead of just a single cross-section in the traditional method. This single cross-section depends on the location, creating uncertainty in the ground truth data. This uncertain data can guide the training process instead of dictating it if the novel hybrid training approach proposed in this work is used. We demonstrate that this hybrid approach performs better than either completely supervised or unsupervised training.

In summary, we have shown that our model architecture and training approach yield high prediction accuracy for the training materials, and equal to lower performance on new alloys. Even though our dataset only contains 4 alloy systems, we observe high prediction accuracy across the different training cases. This highlights the importance of ensuring the training data has a variety of alloy systems that cover the range of possible melt pool morphologies that are likely to be seen during its use.

The study was completed with a printed set of sample cubes to demonstrate that the powder-free single-track morphology is a useful indicator for the quality of the final printed component. We conclude that a neural network trained only on single-track morphologies proves to be a useful tool to anticipate the final part quality.

This work sets a new path for using unsupervised machine learning in parameter optimisation for laser powder bed fusion as a low time and labour cost framework with high prediction and generalisation accuracies. This framework will be of interest to those looking for a faster method to determine the printing parameters for novel materials, or new machine hardware.

References

[1] X. P. Tan, Y. J. Tan, C. S.L. Chow, S. B. Tor, and W. Y. Yeong. Metallic powder-bed based 3d printing of cellular scaffolds for orthopaedic implants: A state-of-the-art review on manufacturing, topological design, mechanical properties and biocompatibility, 7 2017.

[2] Amir M. Mirzendehdel and Krishnan Suresh. Support structure constrained topology optimization for additive manufacturing. *CAD Computer Aided Design*, 81:1–13, 2016.

[3] Alain Garaigordobil, Rubén Ansola, Estrella Veguería, and Igor Fernandez. Overhang constraint for topology optimization of self-supported compliant mechanisms considering additive manufacturing. *CAD Computer Aided Design*, 109:33–48, 4 2019.

[4] Weiming Wang, Dirk Munro, Charlie C. L. Wang, Fred van Keulen, and Jun Wu. Space-time topology optimization for additive manufacturing. *Structural and Multidisciplinary Optimization*, 61:1–18, 2020.

[5] Francesco Tamburrino, Serena Graziosi, and Monica Bordegoni. The design process of additively manufactured mesoscale lattice structures: A review. *Journal of Computing and Information Science in Engineering*, 18, 12 2018.

[6] David Downing, Martin Leary, Matthew McMillan, Ahmad Alghamdi, and Milan Brandt. Heat transfer in lattice structures during metal additive manufacturing: numerical exploration of temperature field evolution. *Rapid Prototyping Journal*, 26:911–928, 5 2020.

[7] A. C. Field, L. N. Carter, N. J.E. Adkins, M. M. Attallah, M. J. Gorley, and M. Strangwood. The effect of powder characteristics on build quality of high-purity tungsten produced via laser powder bed fusion (lpbf). *Metallurgical and Materials Transactions A: Physical Metallurgy and Materials Science*, 51:1367–1378, 3 2020.

[8] A. Engelhardt, M. Kahl, J. Richter, P. Krooß, A. Kroll, and T. Niendorf. Investigation of processing windows in additive manufacturing of alsi10mg for faster production utilizing data-driven modeling. *Additive Manufacturing*, 55, 7 2022.

[9] Holden Hyer, Le Zhou, Sharon Park, Guilherme Gottsfritz, George Benson, Bjorn Tolentino, Brandon McWilliams, Kyu Cho, and Yongho Sohn. Understanding the laser powder bed fusion of alsi10mg alloy. *Metallography, Microstructure, and Analysis*, 9:484–502, 8 2020.
[10] Mireia Vilanova, Rubén Escribano-garcía, Teresa Guraya, and Maria San Sebastian. Optimizing laser powder bed fusion parameters for in-738lc by response surface method. *Materials*, 13:1–12, 11 2020.
[11] Qixiang Luo, Lu Yin, Timothy W. Simpson, and Allison M. Beese. Effect of processing parameters on pore structures, grain features, and mechanical properties in ti-6al-4v by laser powder bed fusion. *Additive Manufacturing*, 56, 8 2022.
[12] W. Xu, E. W. Lui, A. Pateras, M. Qian, and M. Brandt. In situ tailoring microstructure in additively manufactured ti-6al-4v for superior mechanical performance. *Acta Materialia*, 125:390–400, 2 2017.
[13] P. I. Pradeep, V. Anil Kumar, A. Sriranganath, Satish Kumar Singh, Ankit Sahu, T. Sasi Kumar, P. Ramesh Narayanan, M. Arumugam, and M. Mohan. Characterization and qualification of lpbf additively manufactured aisi-316l stainless steel brackets for aerospace application. *Transactions of the Indian National Academy of Engineering*, 5:603–616, 9 2020.
[14] Cassidy Silbernagel, Adedeji Aremu, and Ian Ashcroft. Using machine learning to aid in the parameter optimisation process for metal-based additive manufacturing. *Rapid Prototyping Journal*, 26:625–637, 2019.
[15] Qian Liu, Hongkun Wu, Moses J. Paul, Peidong He, Zhongxiao Peng, Bernd Gludovatz, Jamie J. Kruzic, Chun H. Wang, and Xiaopeng Li. Machine-learning assisted laser powder bed fusion process optimization for alsi10mg: New microstructure description indices and fracture mechanisms. *Acta Materialia*, 201:316–328, 12 2020.
[16] G. Tapia, A. H. Elwany, and H. Sang. Prediction of porosity in metal-based additive manufacturing using spatial gaussian process models. *Additive Manufacturing*, 12:282–290, 10 2016.
[17] Mostafa Yakout, M. A. Elbestawi, and Stephen C. Veldhuis. Density and mechanical properties in selective laser melting of invar 36 and stainless steel 316l. *Journal of Materials Processing Technology*, 266:397–420, 4 2019.
[18] Luke Scime and Jack Beuth. A multi-scale convolutional neural network for autonomous anomaly detection and classification in a laser powder bed fusion additive manufacturing process. *Additive Manufacturing*, 24:273–286, 12 2018.
[19] Luke Scime and Jack Beuth. Melt pool geometry and morphology variability for the inconel 718 alloy in a laser powder bed fusion additive manufacturing process. *Additive Manufacturing*, 29, 10 2019.
[20] Haijun Gong, Hengfeng Gu, Kai Zeng, J J S Dilip, Deepankar Pal, Brent Stucker, Daniel Christiansen, Jack Beuth, and John J Lewandowski. Melt pool characterization for selective laser melting of ti-6al-4v pre-alloyed powder. pages 256–267, 2014.
[21] E. Alabort, Y. T. Tang, D. Barba, and R. C. Reed. Alloys-by-design: A low-modulus titanium alloy for additively manufactured biomedical implants. *Acta Materialia*, 229, 5 2022.
[22] Johannes Trapp, Alexander M. Rubenchik, Gabe Guss, and Manyalibo J. Matthews. In situ absorptivity measurements of metallic powders during laser powder-bed fusion additive manufacturing. *Applied Materials Today*, 9:341–349, 12 2017.
[23] Kurian Antony, N. Arivazhagan, and K. Senthilkumaran. Numerical and experimental investigations on laser melting of stainless steel 316l metal powders. *Journal of Manufacturing Processes*, 16:345–355, 2014.
[24] L. Parry, I. A. Ashcroft, and R. D. Wildman. Understanding the effect of laser scan strategy on residual stress in selective laser melting through thermo-mechanical simulation. *Additive Manufacturing*, 12:1–15, 10 2016.
[25] Mohamed Amine Bouhlel, John T. Hwang, Nathalie Bartoli, Rémi Lafage, Joseph Morlier, and Joaquim R.R.A. Martins. A python surrogate modeling framework with derivatives. *Advances in Engineering Software*, 135:102662, 2019.
[26] Ruichen Jin, Wei Chen, and Agus Sudjianto. An efficient algorithm for constructing optimal design of computer experiments. *Journal of Statistical Planning and Inference*, 134:268–287, 9 2005.
[27] Wayne E. King, Holly D. Barth, Victor M. Castillo, Gilbert F. Gallegos, John W. Gibbs, Douglas E. Hahn, Chandrika Kamath, and Alexander M. Rubenchik. Observation of keyhole-mode laser melting in laser powder-bed fusion additive manufacturing. *Journal of Materials Processing Technology*, 214:2915–2925, 2014.
[28] Christian Tenbrock, Felix Gabriel Fischer, Konrad Wissenbach, Johannes Henrich Schleifenbaum, Philipp Wagenblast, Wilhelm Meiners, and Joachim Wagner. Influence of keyhole and conduction mode melting for top-hat shaped beam profiles in laser powder bed fusion. *Journal of Materials Processing Technology*, 278, 4 2020.
[29] Sagar Patel and Mihaela Vlasea. Melting modes in laser powder bed fusion. *Materialia*, 9, 3 2020.
[30] Johannes Schindelin, Ignacio Arganda-Carreras, Erwin Frise, Verena Kaynig, Mark Longair, Tobias Pietzsch, Stephan Preibisch, Curtis Rueden, Stephan Saalfeld, Benjamin Schmid, Jean-Yves Tinevez, Daniel James White, Volker Hartenstein, Kevin Eliceiri, Pavel Tomancak, and Albert Cardona. Fiji: an open-source platform for biological-image analysis. *Nature Methods*, 9:676–682, 7 2012.
[31] Mohammad Parsazadeh, Shashank Sharma, and Narendra Dahotre. Towards the next generation of machine learning models in additive manufacturing: A review of process dependent material evolution. *Progress in Materials Science*, 135:101102, 6 2023.
[32] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. Imagenet classification with deep convolutional neural networks. *Communications of the ACM*, 60:84–90, 5 2017.
[33] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *CoRR*, abs/1409.1556, 9 2014.
[34] Yann LeCun, Koray Kavukcuoglu, and Clement Farabet. Convolutional networks and applications in vision. pages 253–256. IEEE, 5 2010.
[35] Xu Ji, Andrea Vedaldi, and Joao Henriques. Invariant information clustering for unsupervised image classification and segmentation. volume 2019-October, pages 9864–9873. IEEE, 10 2019.
[36] Erik G Learned-Miller. Entropy and mutual information. *Department of Computer Science, University of Massachusetts, Amherst*, 4, 2013.
[37] Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, Manjunath Kudlur, Josh Levenberg, Rajat Monga, Sherry Moore, Derek G. Murray, Benoit Steiner, Paul Tucker, Vijay Vasudevan, Pete Warden, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. Tensorflow: A system for large-scale machine learning. page 44, 5 2016.
[38] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. 12 2015.
[39] Min Lin, Qiang Chen, and Shuicheng Yan. Network in network. 12 2013.
[40] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. 12 2014.
[41] Manyalibo Matthews, Johannes Trapp, Gabe Guss, and Alexander Rubenchik. Direct measurements of laser absorptivity during metal melt

pool formation associated with powder bed fusion additive manufacturing processes. *Journal of Laser Applications*, 30, 8 2018.

[42] Sagar Patel, Haoxiu Chen, Mihaela Vlasea, and Yu Zou. The influence of beam focus during laser powder bed fusion of a high reflectivity aluminium alloy — alsi10mg. *Additive Manufacturing*, 59, 11 2022.

[43] Meet Gor, Harsh Soni, Vishal Wankhede, Pankaj Sahlot, Krzysztof Grzelak, Ireneusz Szachgluchowicz, and Janusz Kluczyński. A critical review on effect of process parameters on mechanical and microstructural properties of powder-bed fusion additive manufacturing of ss316l, 11 2021.

[44] Liu Cao and Wei Guan. Simulation and analysis of lpbf multi-layer single-track forming process under different particle size distributions. *The International Journal of Advanced Manufacturing Technology*, 114:2141–2157, 6 2021.

[45] Sabrine Ziri, Anis Hor, and Catherine Mabru. Combined effect of powder properties and process parameters on the density of 316l stainless steel obtained by laser powder bed fusion. *International Journal of Advanced Manufacturing Technology*, 120:6187–6204, 6 2022.

[46] Minhalina Ahmad Buhairi, Farhana Mohd Foudzi, Fathin Iliana Jamhari, Abu Bakar Sulong, Nabilah Afiqah Mohd Radzuan, Norhamidi Muhamad, Intan Fadhlina Mohamed, Abdul Hadi Azman, Wan Sharuzi Wan Harun, and M S H Al-Furjan. Review on volumetric energy density: influence on morphology and mechanical properties of ti6al4v manufactured via laser powder bed fusion. *Progress in Additive Manufacturing*, 8:265–283, 4 2023.

Appendix A. Model architecture

Appendix A.1. Residual block

Figure A.15 shows the architecture of the residual block used in the neural network. This block uses the standard residual block design as specified in the original work by He et al. [38]. The block contains 2 pathways, one with 2 3x3 convolutional layers, and a skip connection with a single 1x1 convolutional layer. The main idea behind this architecture is to allow the input information to propagate deeper into a network, and it has been shown that using this network permits deeper networks to be trained with higher accuracy. In this work, it was found that using this residual block in place of an equivalent block without the skip connection improved the generalisation accuracy, so it was used for the final version of the network.

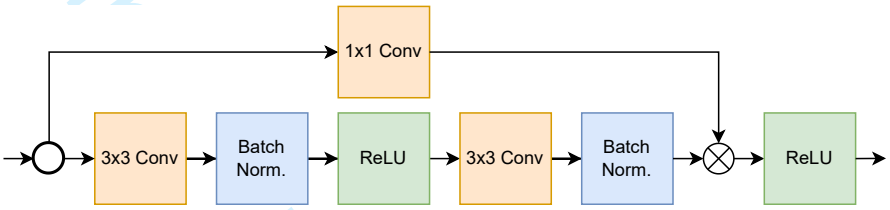

Figure A.15: The architecture of the residual block used in the convolution model. *Authors own work.*

Appendix A.2. Output head

Appendix A.2 shows the architecture of the output heads used in this work. As mentioned in Section 3.2, each output head consists of 2 hidden dense layers of 32 neurons, with a final output layer of length 4, corresponding to the desired number of output clusters. Each of the hidden dense layers is followed by a batch normalisation layer and then a ReLU activation layer, not shown in Appendix A.2 for brevity. What is shown is the softmax activation function that follows the output layer and a graphical representation of the output. This output is a discrete probability distribution, that gives the likelihood of the input image belonging to each category. The decision of the number of layers and the number of neurons in each layer was made using a combination of trial and error, and by taking the training time into account. As the number of layers and neurons increases, the training time and the chance of the model overfitting both increases. The more complex models were also observed to have a lower generalisation accuracy.

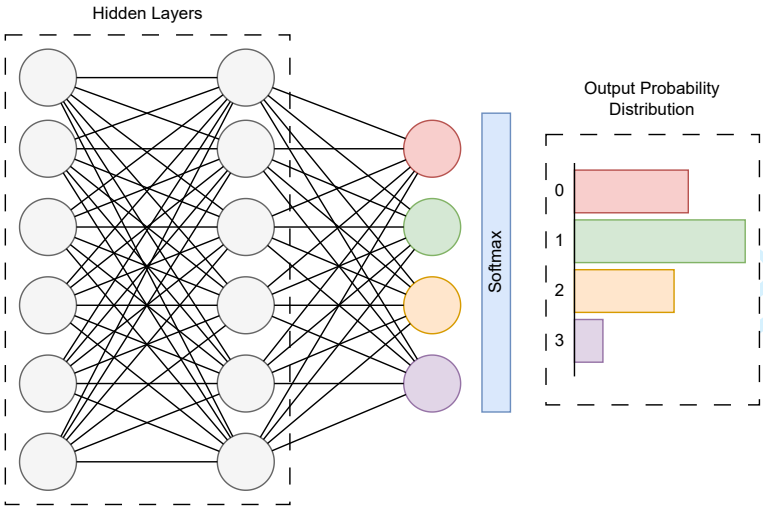

Figure A.16: A diagram of the output head used for this work. For brevity, not all 32 neurons are shown for the dense layers, and the batch normalisation and ReLU activation layers have also been omitted. *Authors own work.*

Appendix B. Data tables

Appendix B.1. Accuracy and error for r ratio variation

B.5 tabulates the data shown in Figure 8. The error for training and generalisation accuracy is calculated by dividing the range of the average accuracy for each by 2. For example, for the case where the r ratio is 0, the highest average generalisation accuracy was for the β-stable titanium alloy at 83.33%, and the lowest was for the Ti-6Al-4V alloy at 33.33%. Half of the range of these two numbers is 25% and is the error quoted.

r Ratio	Training Accuracy (%)	Training Error (%)	Generalisation Accuracy (%)	Generalisation Error (%)
0.00	68.63	19.14	63.12	25.00
0.50	97.65	2.33	80.42	9.17
1.00	93.26	8.45	78.79	10.09

Table B.5: Tabular data showing the training and generalisation accuracies for the 3 different r ratios tested. *Authors own work.*

The DOE for each of the 3 cut alloys, as well as the measured width and depth measurements, are displayed in the following tables, Appendix B.2, Appendix B.3 and Appendix B.4.

Appendix B.2. Ti-6Al-4V

ID	Power (W)	Scan Speed (mm/s)	Width (μm)	Depth (μm)	R	Melting Mode
0	72	1107	65.8	20.7	0.31	Undermelting
1	200	1295	95.4	103	1.08	Keyhole
2	176	965	96.1	121.7	1.27	Keyhole
3	92	1718	54.7	20.7	0.38	Undermelting
4	178	531	118	245.4	2.08	Keyhole
5	143	1176	93.4	73.7	0.79	Optimal
6	131	1377	79.1	50.8	0.64	Optimal
7	128	1755	67.3	39.2	0.58	Balling
8	146	1625	72.5	49	0.68	Optimal
9	195	1048	105.5	133.8	1.27	Keyhole
10	171	1639	81.9	71.5	0.87	Keyhole
11	158	1792	72.3	58.2	0.80	Balling
12	124	641	109.9	139.7	1.27	Keyhole
13	153	1017	107.9	102	0.95	Keyhole
14	108	1866	60.6	23.1	0.38	Balling
15	59	1333	52.7	15.3	0.29	Undermelting
16	136	1984	66	33	0.50	Balling
17	191	775	124.7	197.9	1.59	Keyhole
18	81	1287	63.6	27.1	0.43	Undermelting
19	138	867	110.2	110.9	1.01	Keyhole
20	149	574	139.2	194.4	1.40	Keyhole
21	101	1220	81.1	40.4	0.50	Optimal
22	165	744	126.7	172.3	1.36	Keyhole
23	118	1071	95.1	67.5	0.71	Optimal
24	54	683	63.3	19.5	0.31	Undermelting
25	65	898	58.7	17.3	0.29	Undermelting
26	103	602	96.6	115.1	1.19	Keyhole
27	115	1542	62.6	42.1	0.67	Optimal
28	74	1567	48.1	19	0.40	Undermelting

29	91	917	78.4	52.5	0.67	Optimal
30	111	834	99.1	81.8	0.83	Keyhole
31	169	1173	94.4	98.3	1.04	Keyhole
32	53	1699	43.6	11.3	0.26	Undermelting
33	159	1471	79.8	72.7	0.91	Keyhole
34	182	1412	91.4	90.2	0.99	Keyhole
35	87	1937	49.3	19.2	0.39	Undermelting
36	187	1832	76.4	69	0.90	Balling
37	78	707	88.2	52.5	0.60	Optimal
38	65	1913	41.9	13.4	0.32	Undermelting
39	95	1508	59.1	27.6	0.47	Undermelting

Appendix B.3. Al-Si-10Mg

ID	Power (W)	Scan Speed (mm/s)	Width (μm)	Depth (μm)	R	Melting Mode
0	157	652	79.6	31.1	0.39	Undermelting
1	187	903	140.4	95.1	0.68	Optimal
2	85	1871	38.5	10.9	0.28	Undermelting
3	324	1506	109.3	113.5	1.04	Balling
4	304	1746	110.6	92.5	0.84	Balling
5	151	380	148.4	93.4	0.63	Optimal
6	350	1998	120.5	109	0.90	Balling
7	178	1679	72.7	24.9	0.34	Undermelting
8	101	554	55.1	18.3	0.33	Undermelting
9	133	868	67.1	24.4	0.36	Undermelting
10	172	1257	78.5	27.4	0.35	Undermelting
11	229	724	196.2	175	0.89	Keyhole
12	138	1477	64.3	20.5	0.32	Undermelting
13	196	512	195.9	169.4	0.86	Keyhole
14	319	987	201.1	204.2	1.02	Keyhole
15	103	309	63.1	24.2	0.38	Undermelting
16	199	233	248.9	259.5	1.04	Keyhole
17	381	1057	209	257.4	1.23	Keyhole
18	120	1117	59.6	20.2	0.34	Undermelting
19	268	940	180.9	172.5	0.95	Keyhole
20	128	1794	54.2	18	0.33	Undermelting
21	302	286	326.3	494.1	1.51	Keyhole
22	399	1389	140	159.7	1.14	Balling
23	390	615	263.9	412.8	1.56	Keyhole
24	246	429	254.8	251	0.99	Keyhole
25	211	1459	80.6	33.3	0.41	Undermelting
26	56	1576	36.5	8.1	0.22	Undermelting
27	282	601	244	271.6	1.11	Keyhole
28	63	1040	40.6	8.6	0.21	Undermelting
29	253	1620	116.8	63.6	0.54	Optimal
30	270	1932	104.2	49	0.47	Undermelting
31	287	1316	184.3	132.2	0.72	Optimal
32	72	761	34.8	12.3	0.35	Undermelting
33	361	1724	132.2	136.5	1.03	Balling
34	337	1218	195.7	189.8	0.97	Keyhole
35	220	1841	77.1	33.8	0.44	Undermelting
36	241	1154	169	107.9	0.64	Optimal
37	347	413	312.7	411.2	1.31	Keyhole
38	92	1354	46.7	13.8	0.30	Undermelting
39	365	809	233.1	297.8	1.28	Keyhole

Appendix B.4. 316L Stainless Steel

ID	Power (W)	Scan Speed (mm/s)	Width (μm)	Depth (μm)	R	Melting Mode
0	139	651	133.6	94.0	0.70	Optimal
1	196	615	133.9	163.2	1.22	Keyhole
2	153	840	114.1	76.9	0.67	Optimal
3	182	1166	94.6	76.6	0.81	Keyhole
4	84	577	92.2	47.3	0.51	Optimal
5	192	886	102.6	99.6	0.97	Keyhole
6	116	740	101.7	53.5	0.53	Optimal
7	175	1347	79.5	65.0	0.82	Balling
8	90	822	74.8	31.3	0.42	Undermelting
9	133	921	91.7	57.1	0.62	Optimal
10	186	1759	77.8	56.2	0.72	Balling
11	143	1949	62.4	42.3	0.68	Balling
12	74	1237	50.6	18.6	0.37	Undermelting
13	62	1911	38.1	12.4	0.33	Undermelting
14	156	1491	72.7	64.5	0.89	Balling
15	100	1656	49.7	21.3	0.43	Undermelting
16	68	788	60.9	23.4	0.38	Undermelting
17	93	1097	61.8	23.7	0.38	Undermelting
18	113	1026	78.6	42.9	0.55	Optimal
19	98	1853	50.3	20.4	0.41	Undermelting
20	128	1188	74.2	46.7	0.63	Optimal
21	145	1695	61.5	55.3	0.90	Balling
22	178	724	129.2	128.9	1.00	Keyhole
23	55	1111	44.9	16.9	0.38	Undermelting
24	198	1543	78.1	71.6	0.92	Balling
25	132	1427	66.8	47.6	0.71	Balling
26	107	1289	61.5	28.7	0.47	Undermelting
27	124	1735	54.4	32.8	0.60	Balling
28	59	1597	37.8	14.8	0.39	Undermelting
29	162	522	142.8	178.3	1.25	Keyhole
30	82	1471	47.9	18.9	0.39	Undermelting
31	173	1570	76.9	58.8	0.76	Balling
32	164	1827	69.8	60.6	0.87	Balling
33	150	1257	71	54.1	0.76	Optimal
34	169	975	92.2	85.4	0.93	Keyhole
35	79	1780	43.5	18.3	0.42	Undermelting
36	103	561	120.3	78.1	0.65	Optimal
37	51	1380	39.1	11.5	0.29	Undermelting
38	72	989	55.6	20.1	0.36	Undermelting
39	120	1978	53.5	30.2	0.56	Balling

Appendix B.5. β-Stable Titanium Alloy

ID	Power (W)	Scan Speed (mm/s)	Width (μm)	Depth (μm)	R	Melting Mode
0	247	1710	103.01	75.16	0.73	Balling
1	185	1222	100.54	80.35	0.80	Optimal
2	207	1609	90.52	69.74	0.77	Balling
3	53	2907	35.32	11.58	0.33	Undermelting
4	151	623	151.64	120.01	0.79	Optimal
5	117	2847	58.32	17.5	0.30	Balling
6	89	2716	50.27	16.02	0.32	Undermelting
7	103	2373	53.89	21.44	0.40	Undermelting
8	297	1132	110.73	159.69	1.44	Keyhole
9	139	2137	65.55	28.84	0.44	Balling
10	177	2207	85.26	43.13	0.51	Balling
11	268	804	154.26	197.89	1.28	Keyhole
12	277	3191	100.71	52.25	0.52	Balling
13	120	919	99.23	60.38	0.61	Optimal
14	93	787	80.83	40.17	0.50	Undermelting
15	210	3413	84.28	33.51	0.40	Balling
16	230	944	138.16	139.73	1.01	Keyhole
17	132	2554	61.11	25.38	0.42	Balling
18	148	1667	80.01	44.85	0.56	Balling
19	258	1201	107.94	127.65	1.18	Keyhole
20	57	1523	47.81	18.49	0.39	Undermelting
21	111	3272	53.39	18.49	0.35	Balling
22	65	2249	44.52	18.49	0.42	Undermelting
23	158	1024	115	81.33	0.71	Optimal
24	265	2005	94.46	81.32	0.86	Balling
25	169	3372	72.61	35.24	0.49	Balling
26	291	509	187.61	364.74	1.94	Keyhole
27	128	1377	77.21	43.37	0.56	Optimal
28	252	2451	97.09	60.62	0.62	Balling
29	194	698	148.35	160.67	1.08	Keyhole
30	163	2657	75.74	33.76	0.45	Balling
31	289	2735	104.32	59.88	0.57	Balling
32	75	1100	65.88	27.11	0.41	Undermelting
33	217	2815	90.19	45.85	0.51	Balling
34	241	596	168.09	255.79	1.52	Keyhole
35	84	3093	44.85	13.06	0.29	Undermelting
36	171	1825	81.32	51.26	0.63	Balling
37	78	1792	55.2	14.54	0.26	Undermelting
38	200	1885	89.7	61.62	0.69	Balling
39	225	2069	93.64	61.62	0.66	Balling
40	145	3053	62.76	22.18	0.35	Balling
41	273	1488	99.72	112.87	1.13	Balling
42	235	3224	91.18	44.85	0.49	Balling
43	67	3451	37.46	9.62	0.26	Undermelting
44	283	2314	99.06	73.19	0.74	Balling
45	229	1317	137.51	99.56	0.72	Optimal

46	197	2507	86.74	47.31	0.55	Balling
47	108	1953	60.13	24.16	0.40	Undermelting
48	188	3011	81.81	34.01	0.42	Balling
49	95	1445	66.04	25.89	0.39	Undermelting

Acknowledgements

This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement number 956401.

Grant EQC2019-006491-P funded by MI- CIU/AEI/10.13039/501100011033 and by ERDF A way of making Europe.

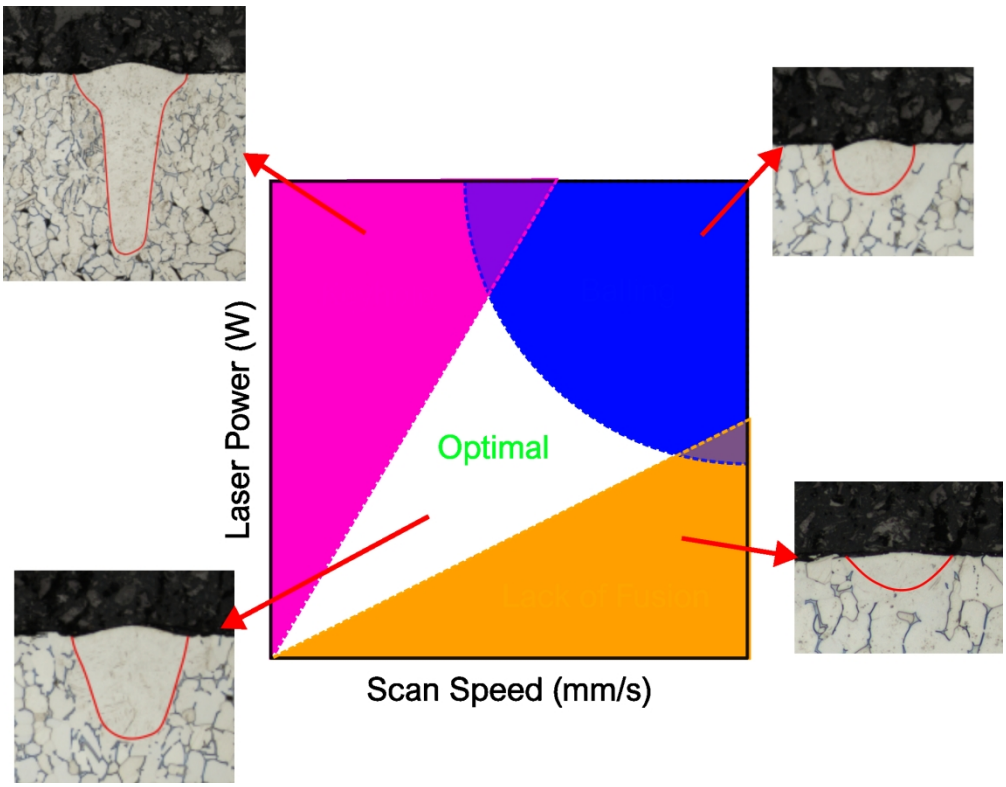

The typical processing map for metals processed using LPBF with the usual melt track morphologies for Ti-6Al-4V indicated. Straight dashed lines on this plot have constant energy density (except the blue balling boundary).

440x342mm (118 x 118 DPI)

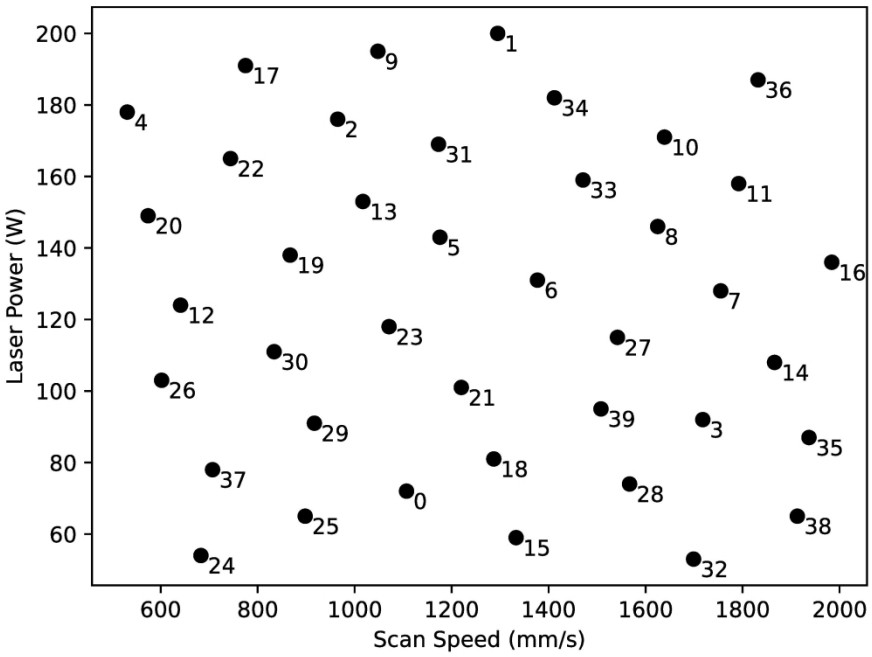

The design of experiments for the Ti-6Al-4V alloy single tracks.

162x121mm (600 x 600 DPI)

Al-Si-10Mg

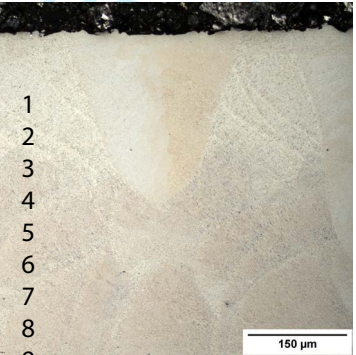

Ti-6Al-4V

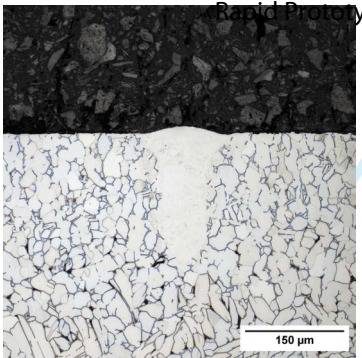

SS316L

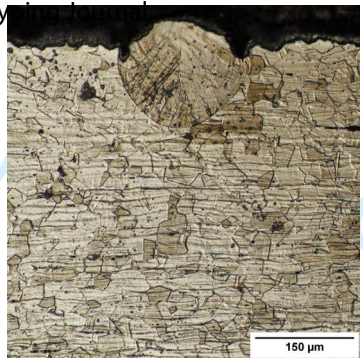

ß-Titanium

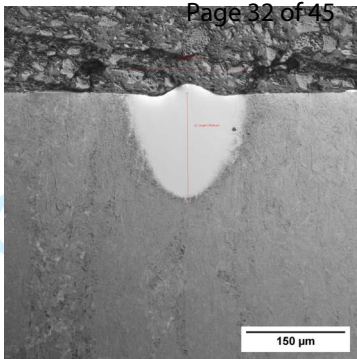

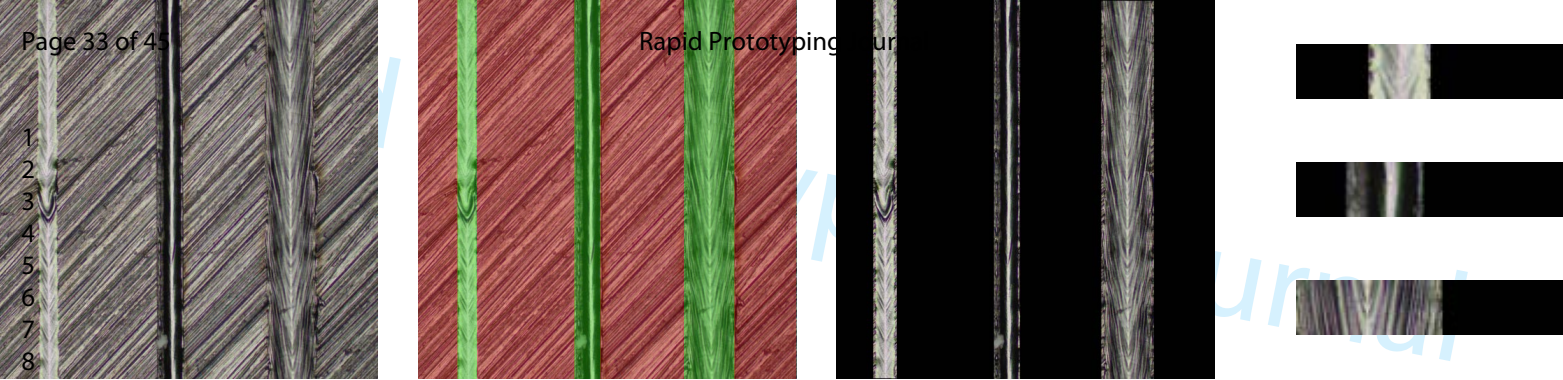
A: Raw Image
B: Segmented Image
C: Background Removed
D: Individual Slices

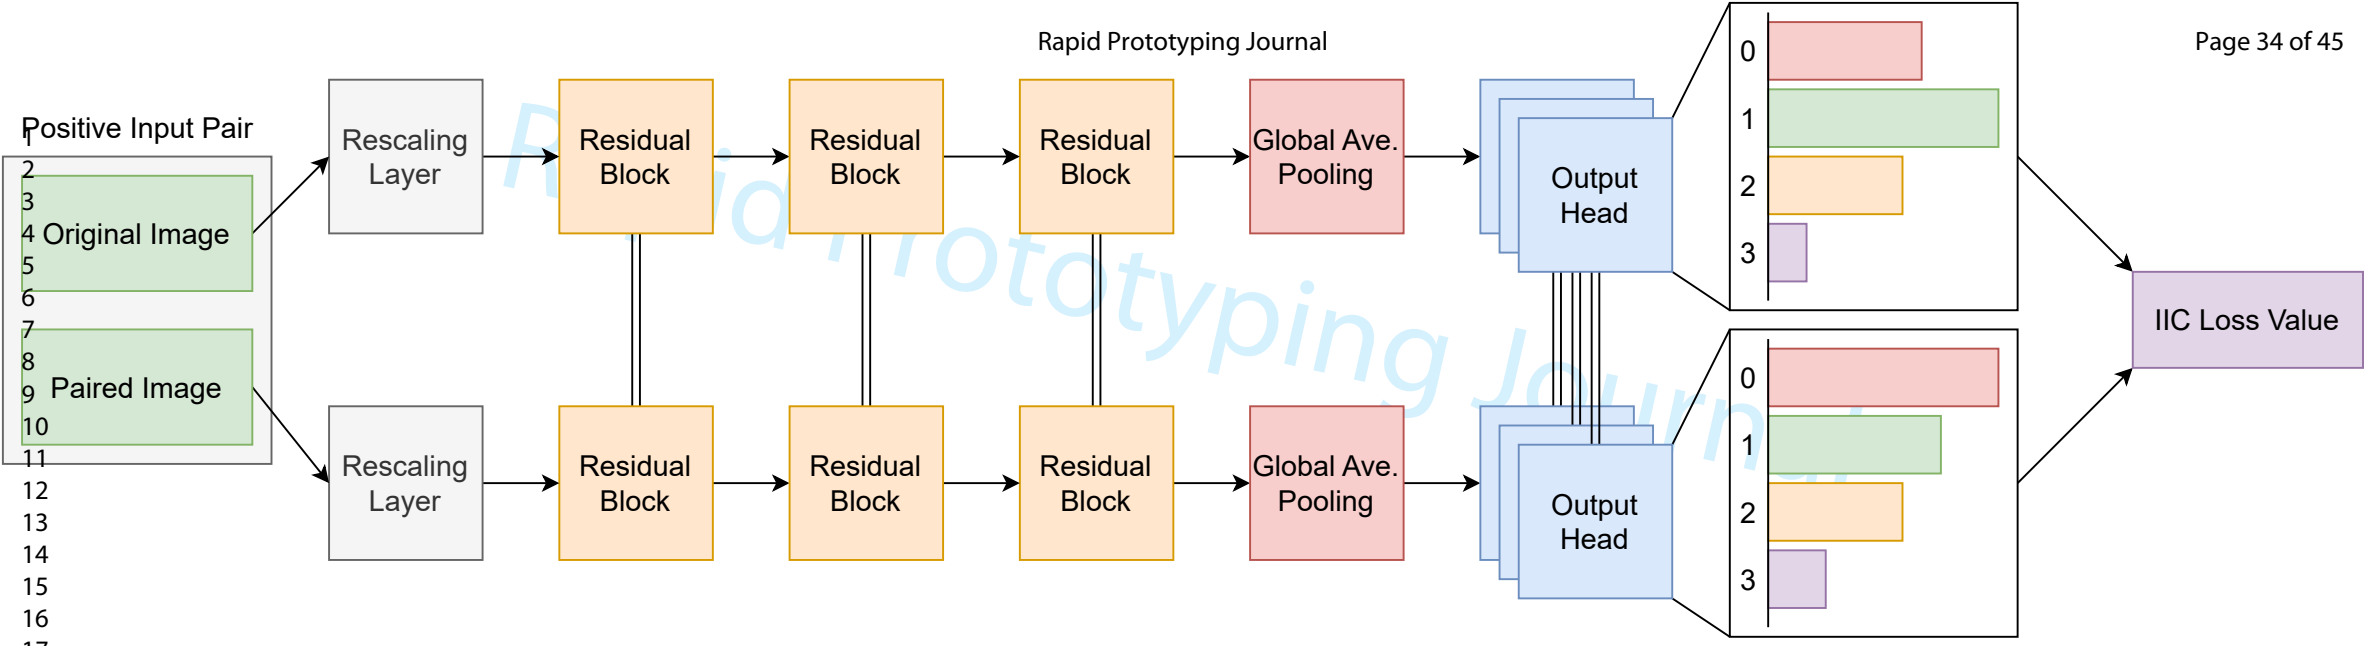

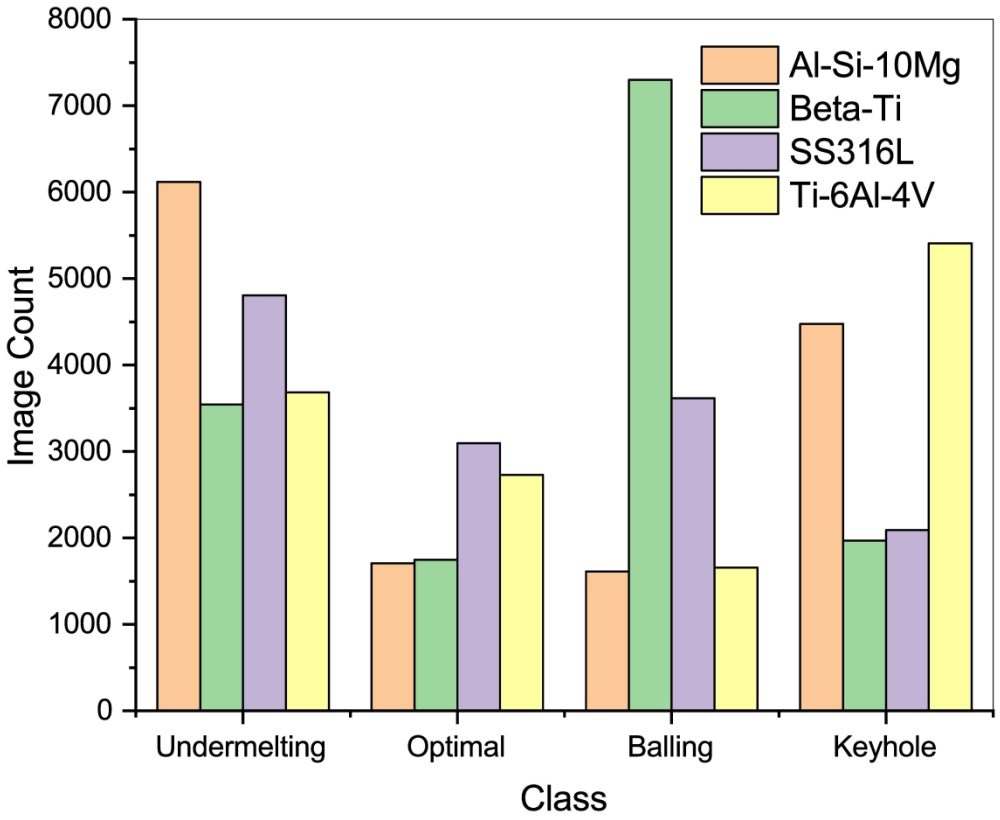

The distribution of classes in the 4 alloys.

216x178mm (600 x 600 DPI)

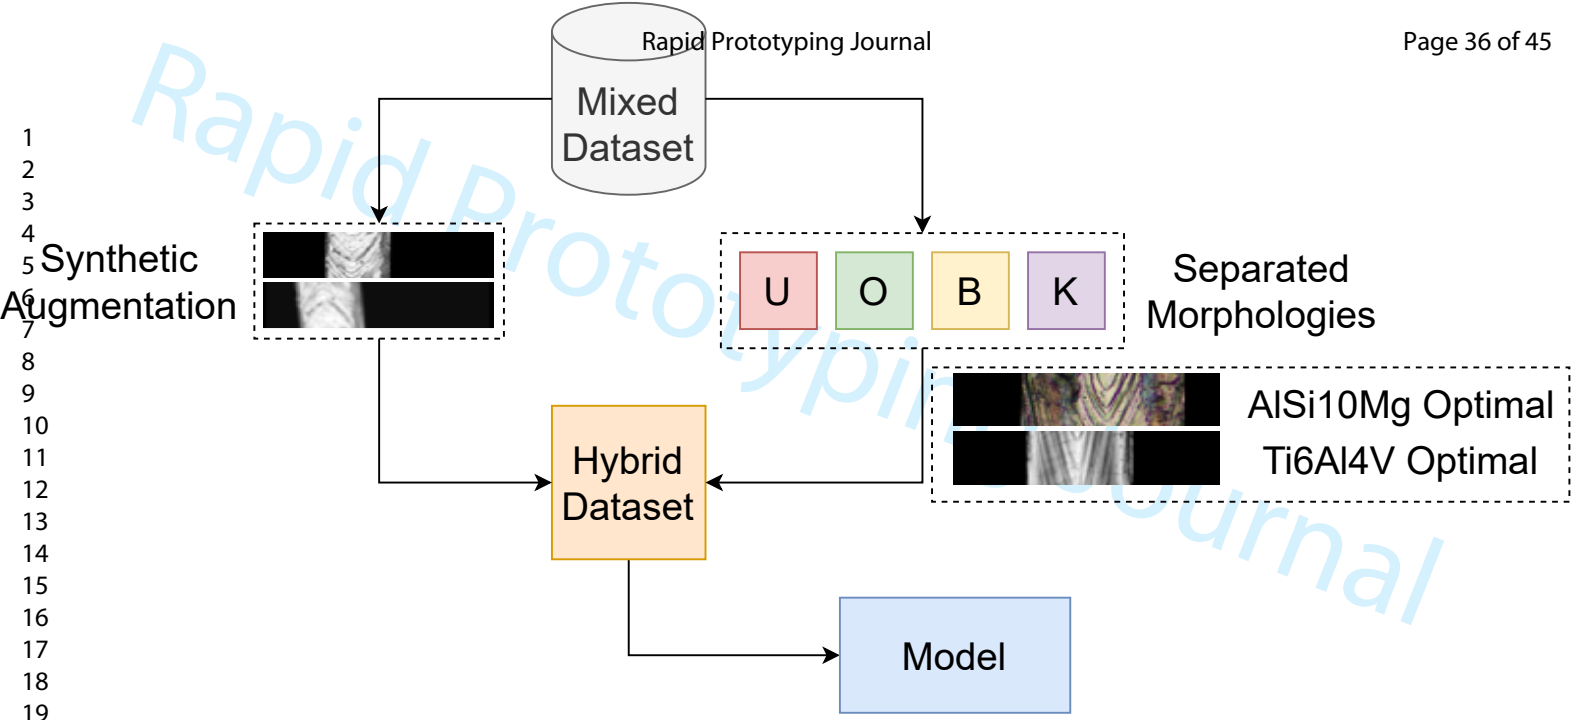

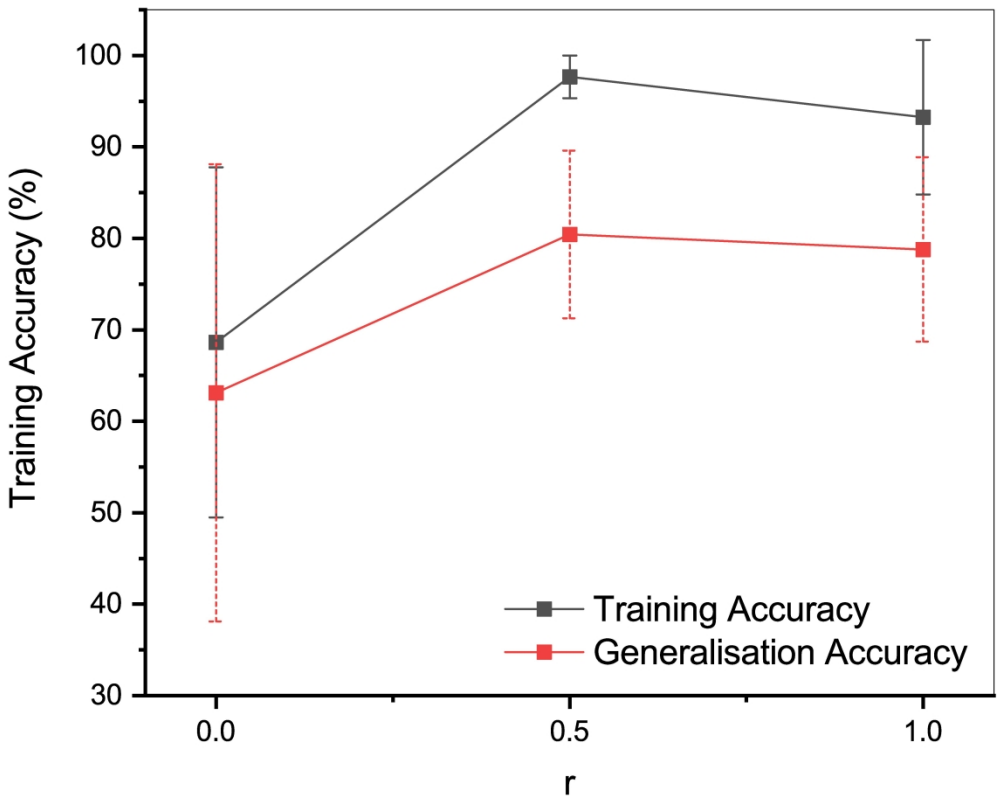

Average training and generalisation accuracy vs ratio r, the fraction of labelled data. The error bars show the range of accuracy for each value of r.

218x174mm (600 x 600 DPI)

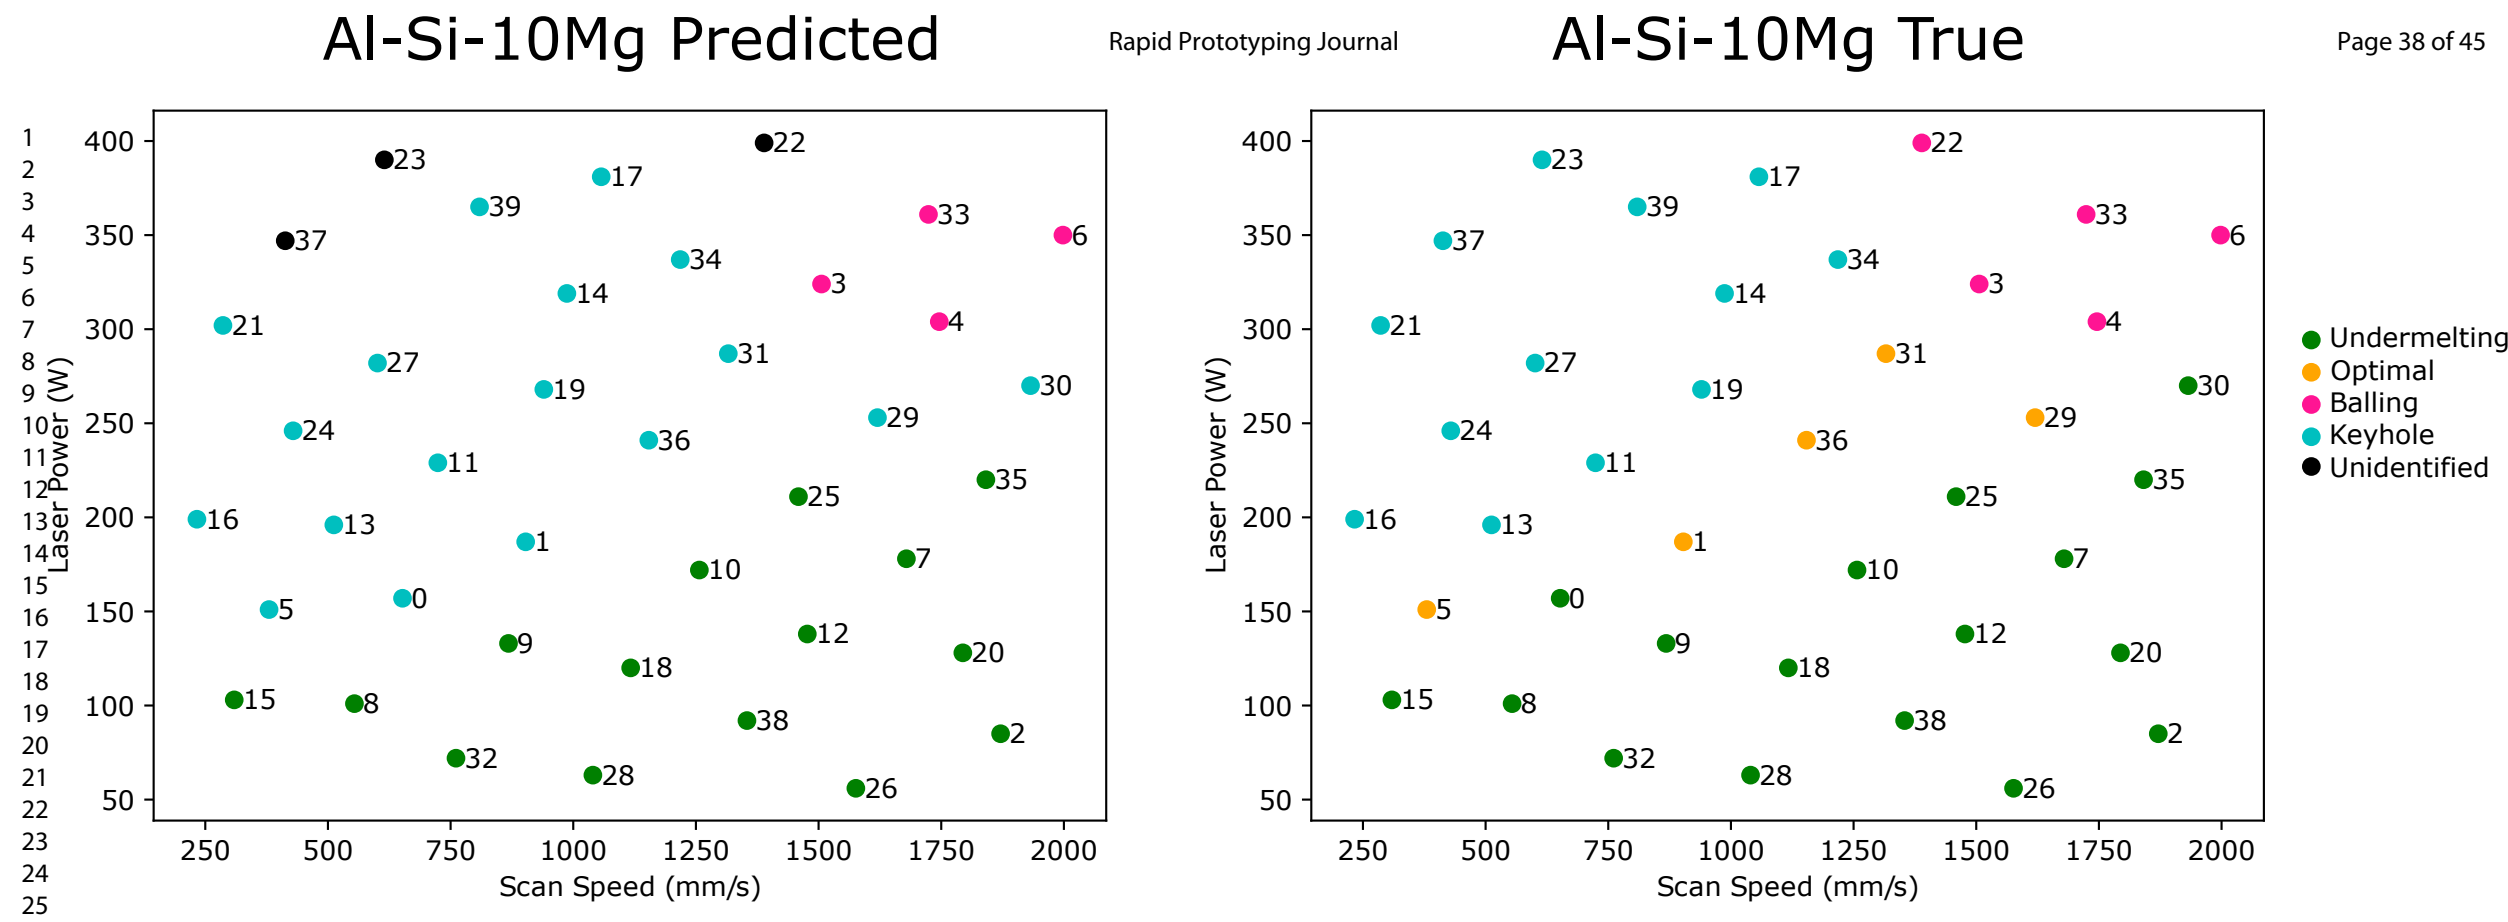
Al-Si-10Mg Predicted
Al-Si-10Mg True
Laser Power (W)
Scan Speed (mm/s)
Undermelting
Optimal
Balling
Keyhole
Unidentified

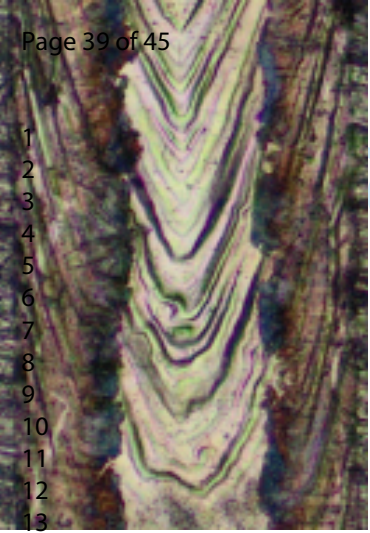

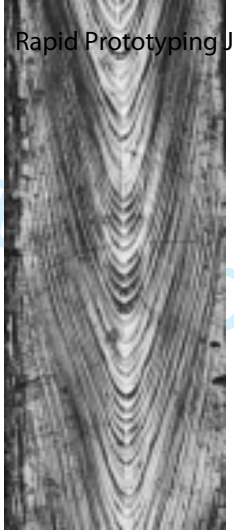

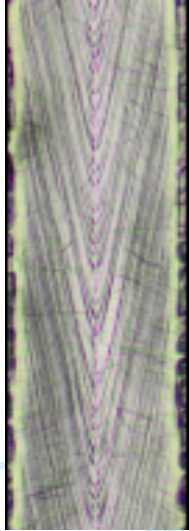

Al-Si-10Mg - 36 ß-titanium - 23 SS316L - 9 Ti-6Al-4V - 6

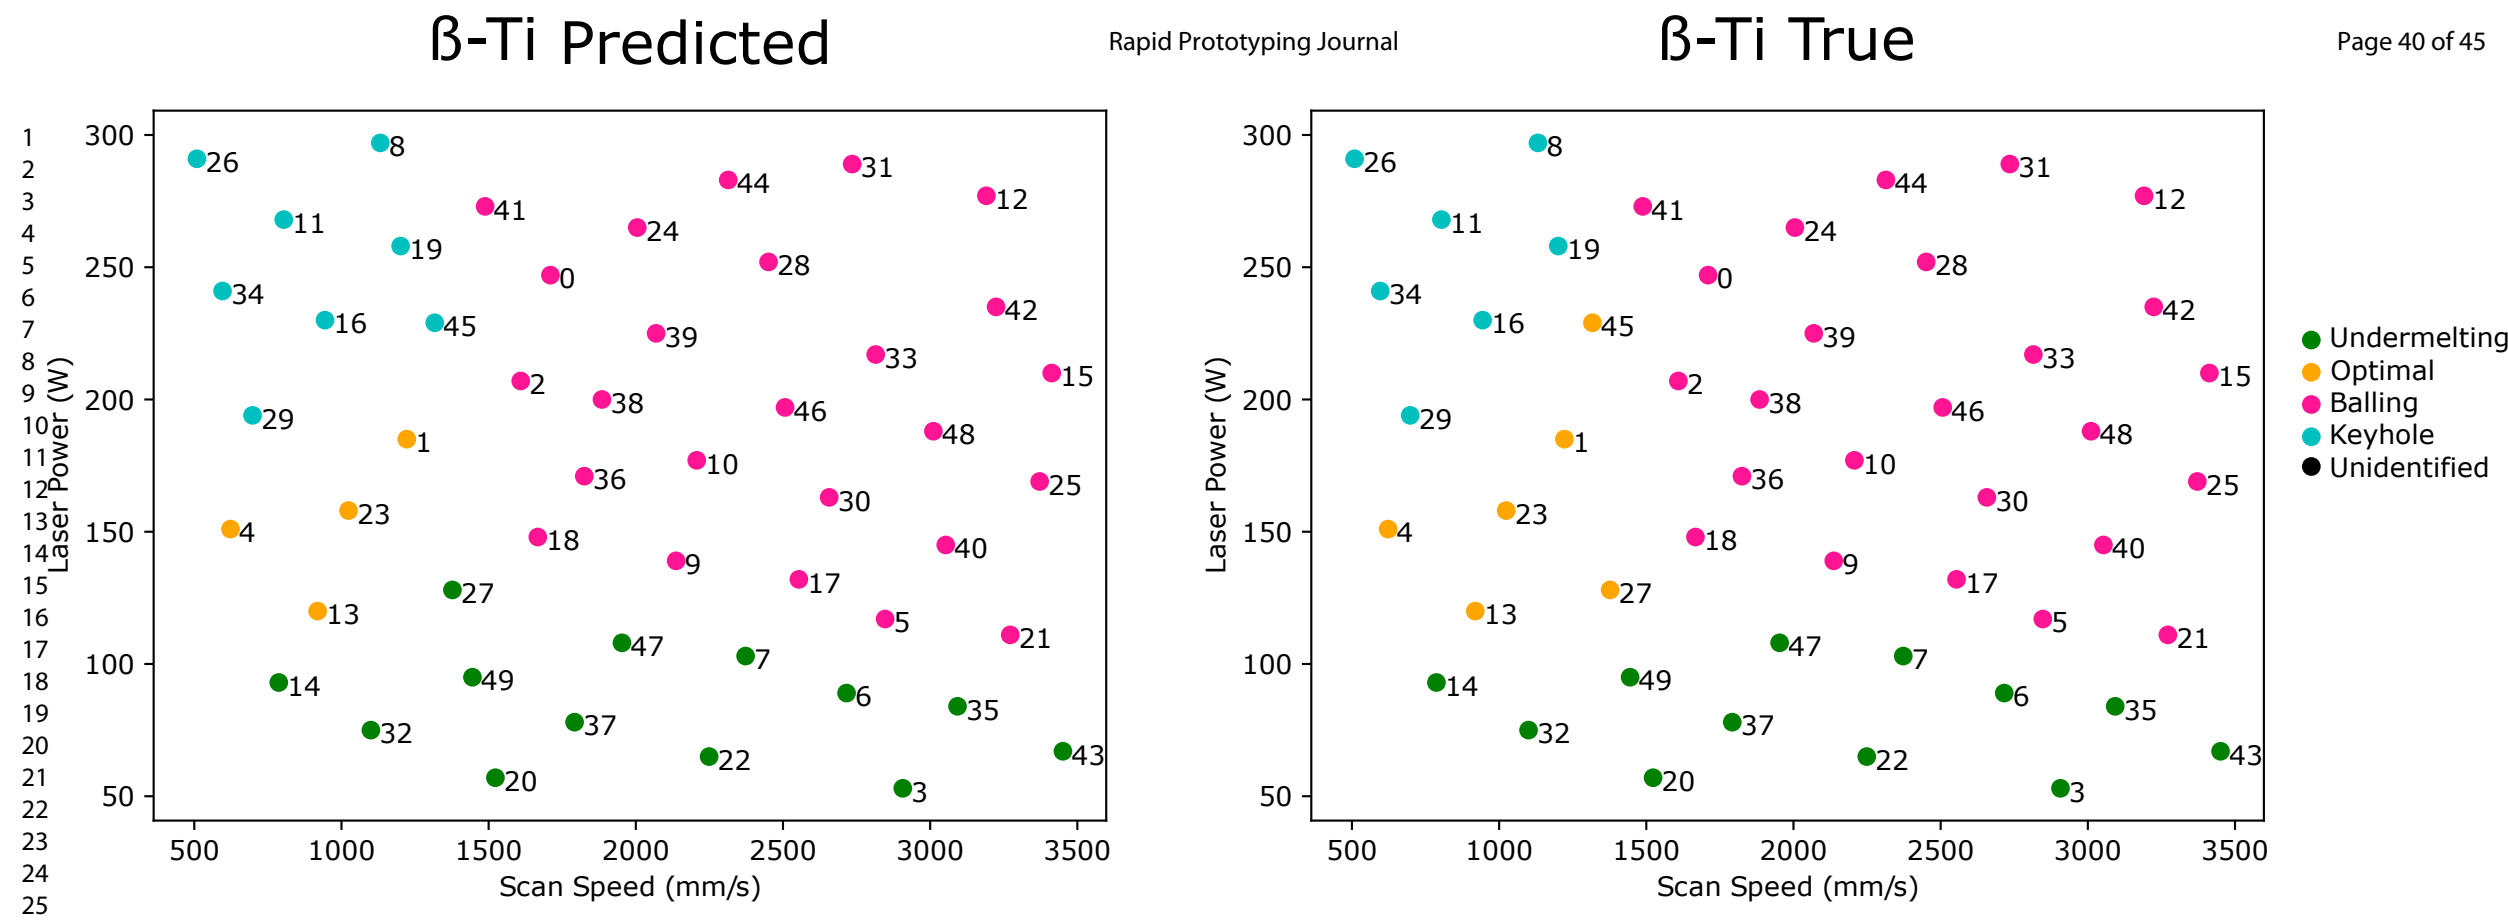
ß-Ti Predicted
ß-Ti True
Laser Power (W)
Scan Speed (mm/s)
300
250
200
150
100
50
500
1000
1500
2000
2500
3000
3500
Undermelting
Optimal
Balling
Keyhole
Unidentified

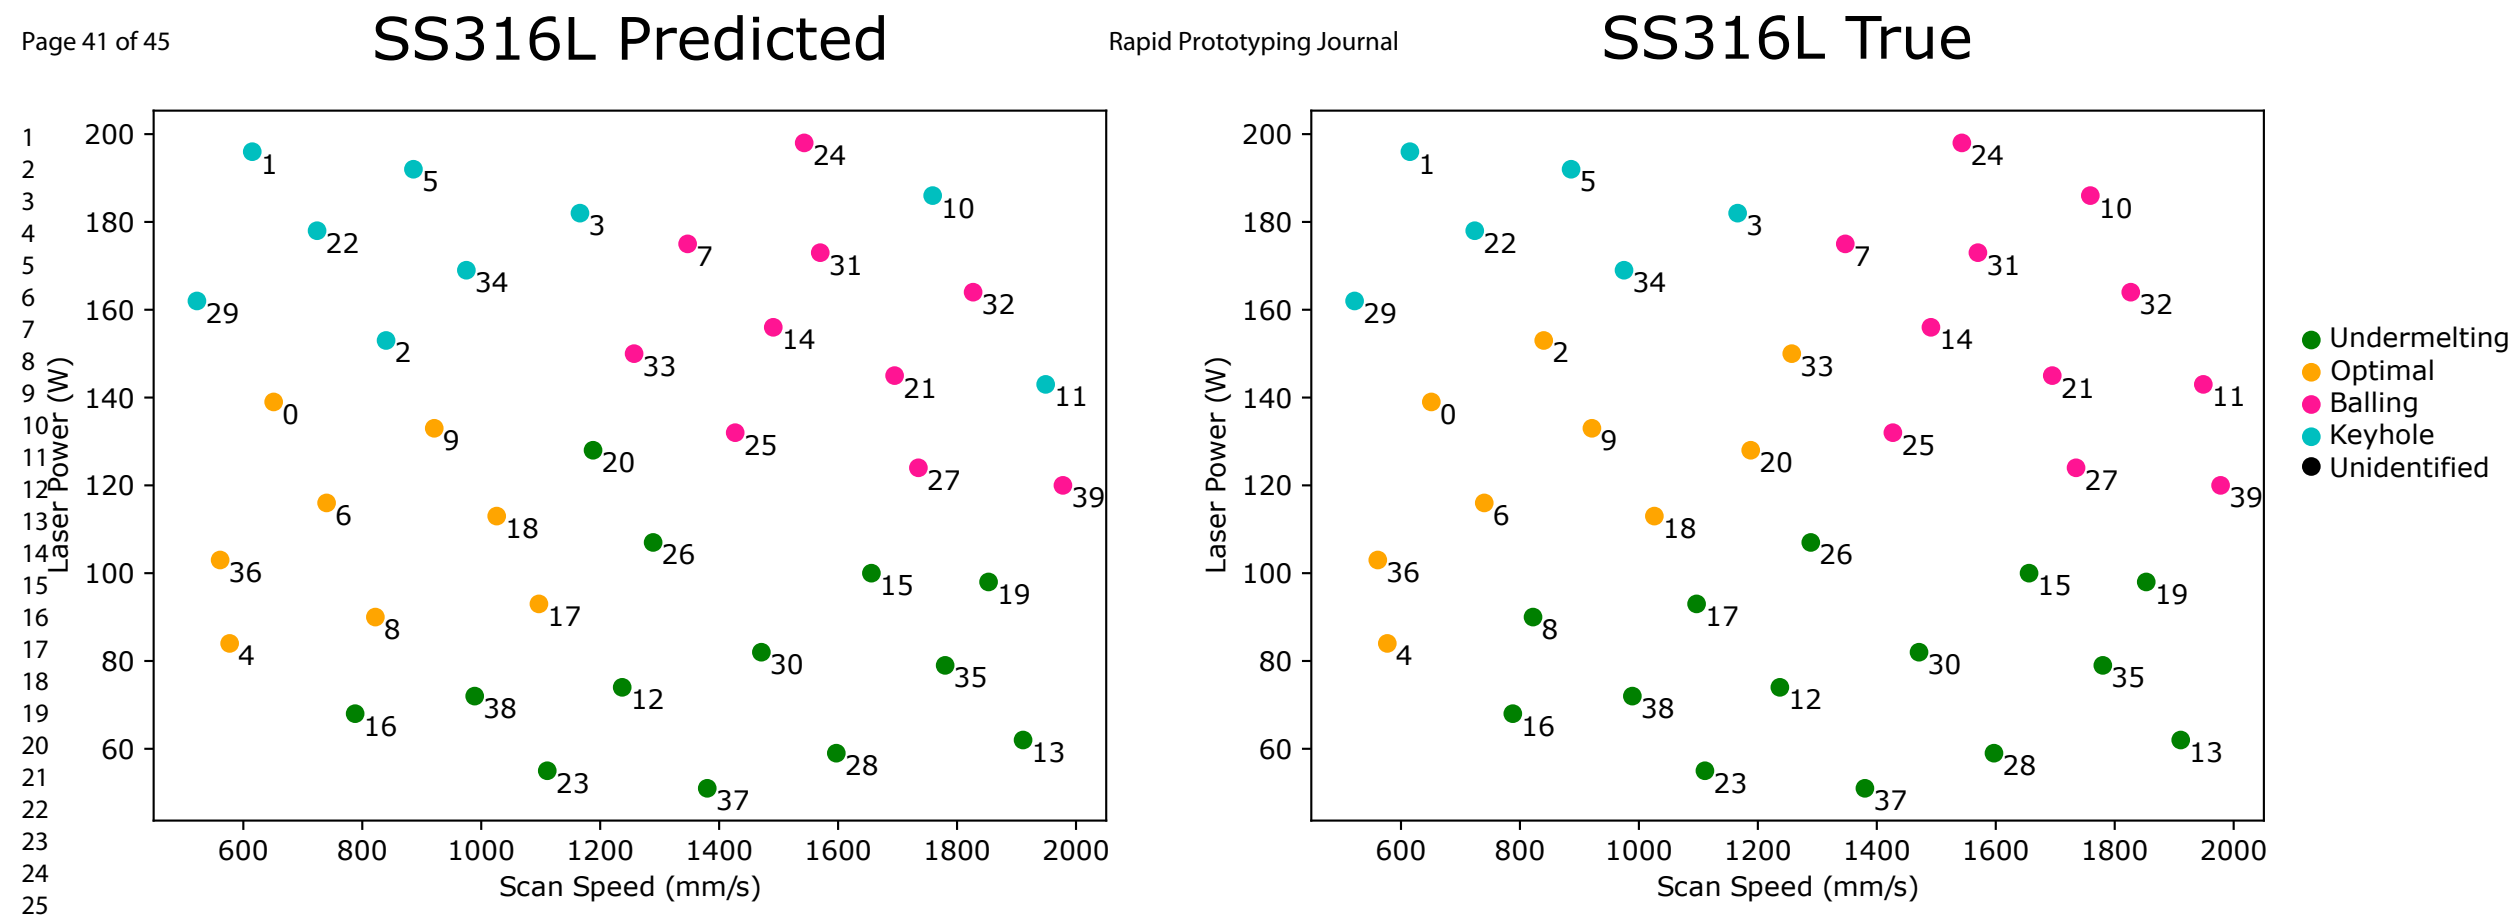
SS316L Predicted
SS316L True
Laser Power (W)
Scan Speed (mm/s)
Undermelting
Optimal
Balling
Keyhole
Unidentified

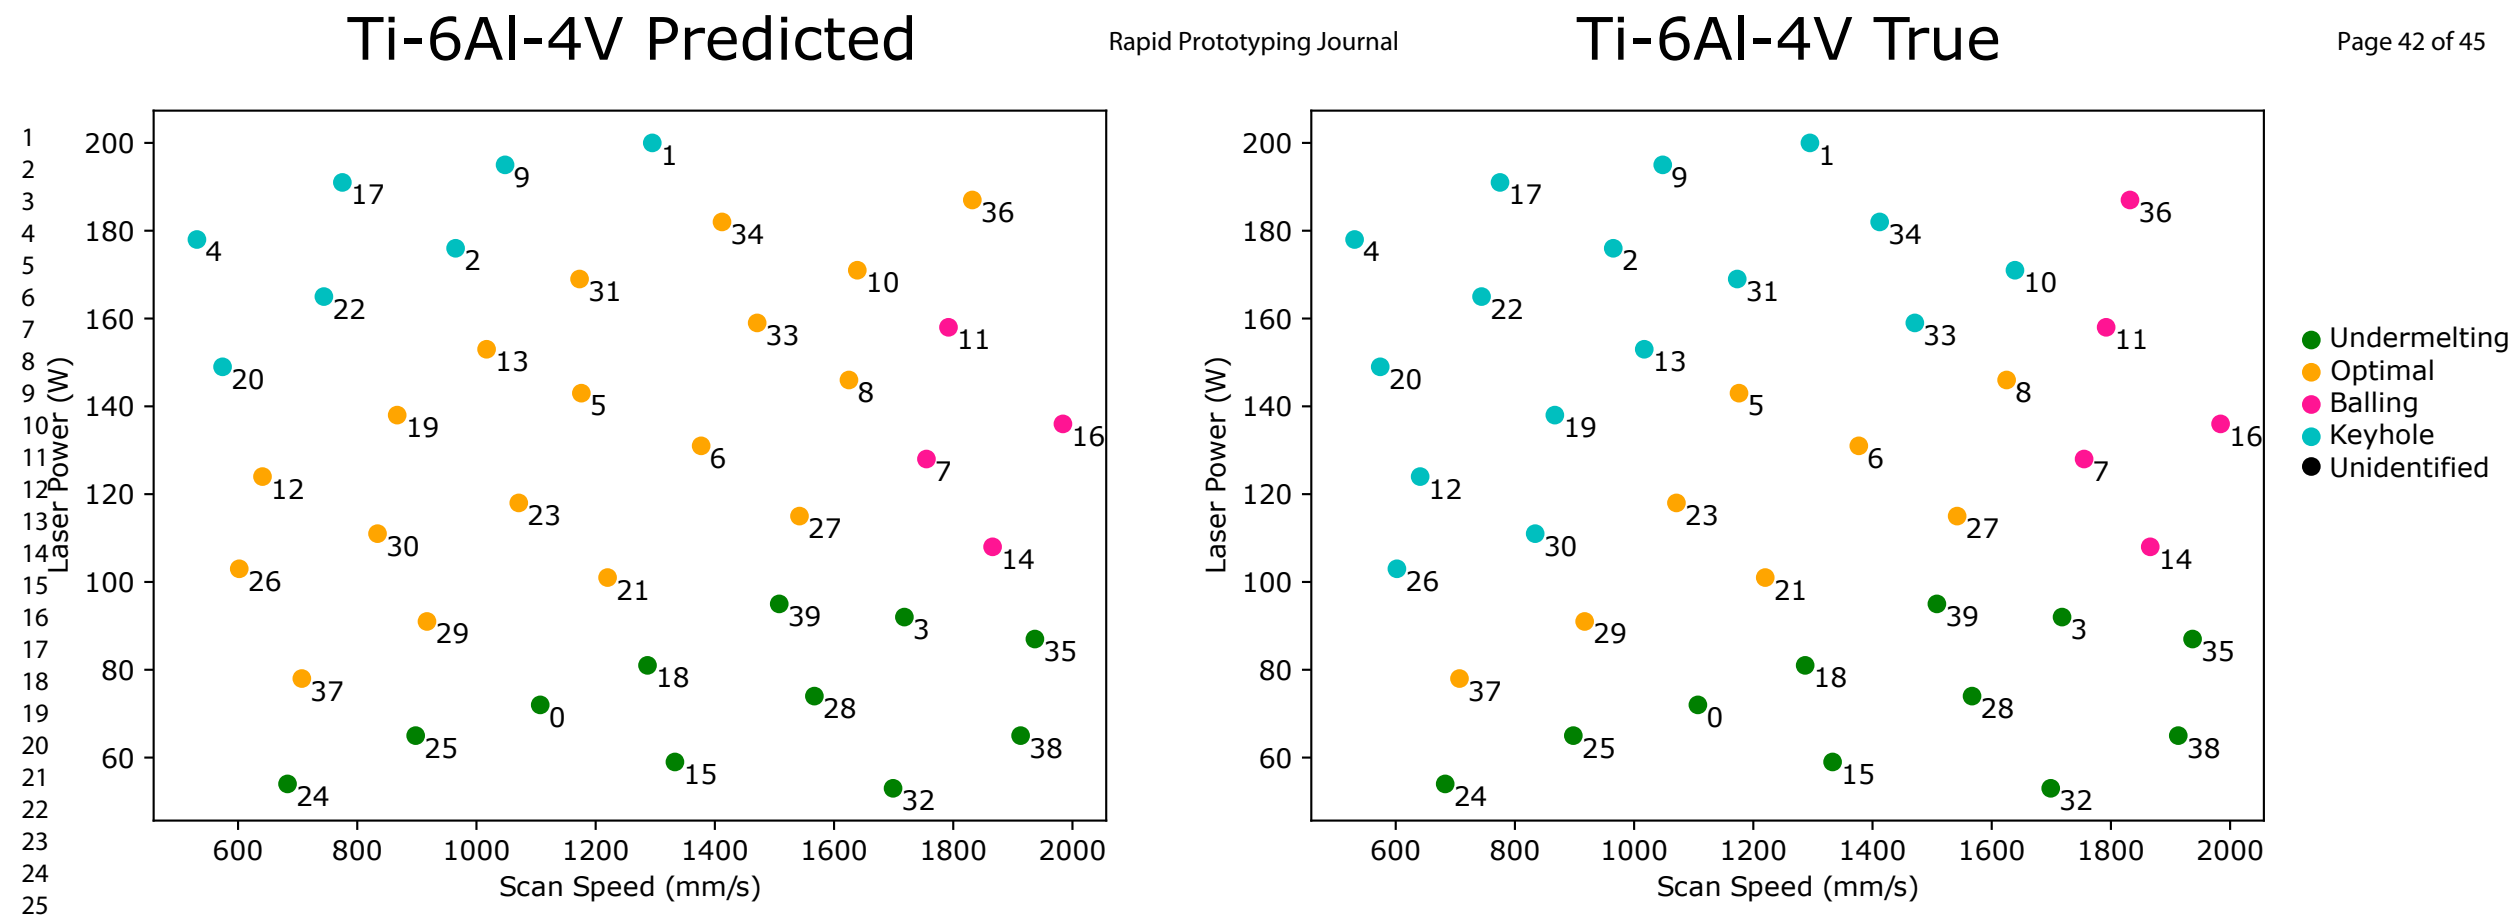
Ti-6Al-4V Predicted
Ti-6Al-4V True
Laser Power (W)
Scan Speed (mm/s)
Undermelting
Optimal
Balling
Keyhole
Unidentified

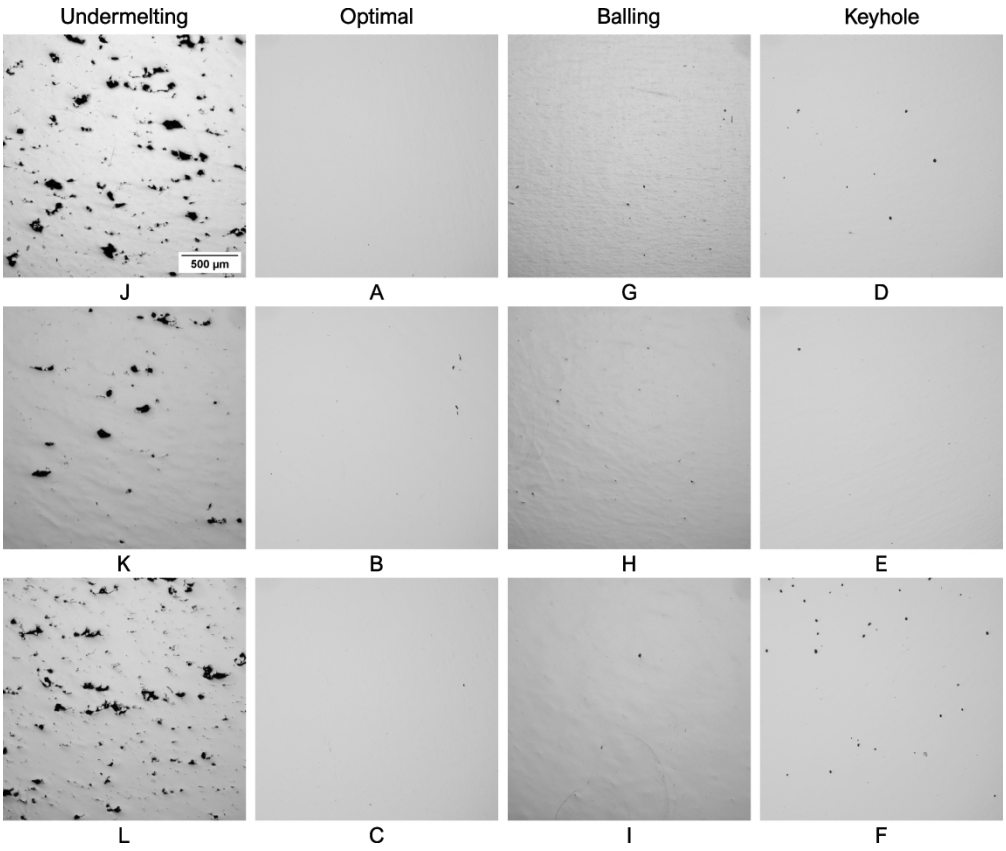

Cross-section images from the cube samples printed in the different melting morphologies as predicted by the ML model. The printing direction is from bottom to top for all images. The images are taken from the samples printed with the parameters in Table 4. All images have the same scale

2112x1761mm (39 x 39 DPI)

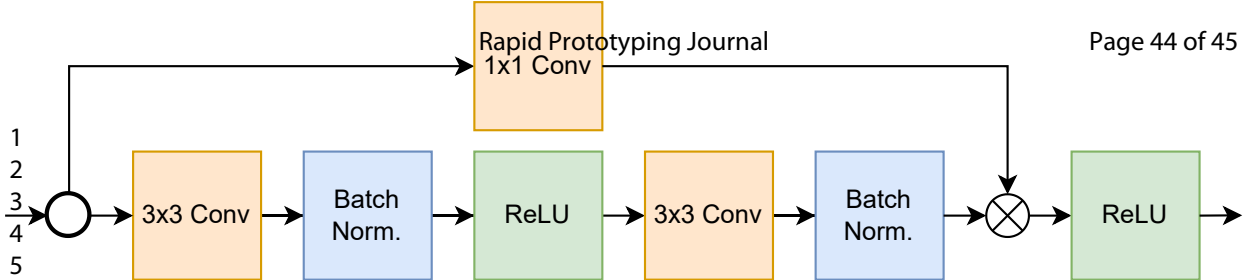
1x1 Conv
1
2
3
4
5
3x3 Conv
Batch Norm.
ReLU
3x3 Conv
Batch Norm.
ReLU

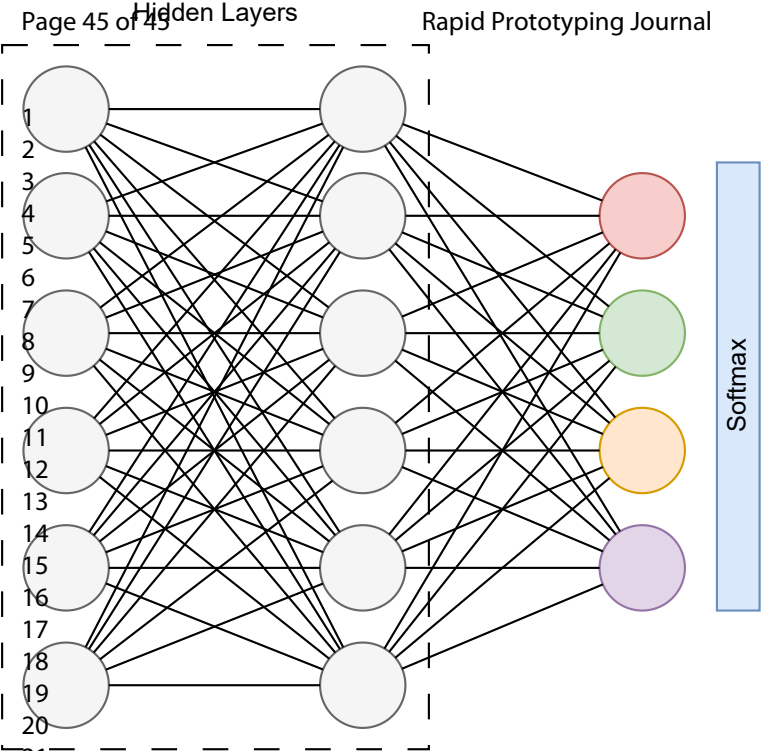
Hidden Layers
1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
Softmax
Output Probability Distribution
0
1
2
3