

Combining High-Throughput Imaging in Visible and SWIR wavelengths for In-Situ Porosity Prediction in Laser Powder Bed Fusion

Ayyoub Ahar^a, Mathieu Vandecasteele^b, Brian G. Booth^b, Kurt De Grave^a, Dries Verhees^a,
Wilfried Philips^b, and Abdellatif Bey-Temsamani^a

^aFlanders Make, 3920 Lommel, Belgium

^bimec TELIN-IPI, Ghent University, 3000 Leuven, Belgium

ABSTRACT

Laser powder bed fusion is at the forefront of manufacturing metallic objects, particularly those with complex geometries or those produced in limited quantities. However, this 3D printing method is susceptible to several printing defects due to the complexities of using a high-power laser with ultra-fast actuation. Accurate online print defect detection is therefore in high demand, and this defect detection must maintain a low computational profile to enable low-latency process intervention. In this work, we propose a low-latency LPBF defect detection algorithm based on fusion of images from high-speed cameras in the visible and short-wave infrared (SWIR) spectrum ranges. First, we design an experiment to print an object while both imposing porosity defects on the print, and recording the laser's melt pool with the high-speed cameras. We then train variational autoencoders on images from both cameras to extract and fuse two sets of corresponding features. The melt pool recordings are then annotated with pore densities extracted from the printed object's CT scan. These annotations are then used to train and evaluate the ability of a fast neural network model to predict the occurrence of porosity from the fused features. We compare the prediction performance of our sensor fused model with models trained on image features from each camera separately. We observe that the SWIR imaging is sensitive to keyhole porosity while the visible-range optical camera is sensitive to lack-of-fusion porosity. By fusing features from both cameras, we are able to accurately predict both pore types, thus outperforming both single camera systems.

Keywords: LPBF, In-situ Monitoring, Lack-of-Fusion, Keyhole Porosity, Additive Manufacturing, SWIR, Variational Autoencoder, Sensor Fusion

1. INTRODUCTION

Additive manufacturing (AM) is a term that encompasses various technologies that create physical objects from digital models by adding material layer by layer.¹ AM offers several advantages over conventional manufacturing methods, such as design flexibility, reduced material waste, and customized production.² Among the different AM techniques, laser powder bed fusion (LPBF) is one of the most widely used and studied.³ LPBF involves melting and solidifying metal powder particles with a laser beam to form solid parts with complex geometries.⁴ LPBF can produce high-quality parts with high mechanical strength, dimensional accuracy, and surface finish, thus making them desirable for applications in aerospace, biomedical, and automotive industries.⁵

However, LPBF also faces several challenges that limit its industrial application, such as high residual stresses, distortion, cracking, and lack of repeatability.⁶ One of the main factors that affect the quality and reliability of LPBF parts is the presence of printing defects, such as porosity, balling, and spatter. These defects can compromise the mechanical performance, fatigue life, and corrosion resistance of the parts, as well as increase the post-processing costs and time.⁷ Therefore, understanding the formation mechanisms and influencing factors of these defects is crucial for improving the LPBF process and achieving defect-free parts.

Porosity in LPBF can be classified into two types: lack-of-fusion porosity and keyhole porosity. Lack-of-fusion porosity occurs when the laser energy input is insufficient to completely melt the powder particles, resulting in

Further author information: (Send correspondence to Ayyoub Ahar)
E-mail: ayyoub.ahar@flandersmake.be

gaps or voids between them.⁸ This type of porosity can be reduced by increasing the laser power, decreasing the scanning speed, or decreasing the hatch spacing, thereby enhancing the melting and coalescence of the powder particles. However, excessive energy input can also lead to another type of porosity: keyhole porosity. Keyhole porosity occurs when the laser energy input is so high that it creates a deep and narrow cavity in the molten pool, a cavity known as and shaped like a keyhole.⁹ The keyhole is filled with vaporized metal and gas, which can be trapped inside it when the molten metal solidifies rapidly. Keyhole porosity can be influenced by various process parameters, such as laser power, scanning speed, hatch spacing, layer thickness, and powder size.¹⁰ Keyhole porosity can also be affected by the material properties, such as melting temperature, thermal conductivity, surface tension, and vapor pressure.

To predict and control for keyhole porosity in LPBF, several analytical and numerical models have been developed in the literature. For example, Wang et al.¹¹ proposed an analytical model, with a closed-form solution, to calculate the keyhole porosity based on a two-dimensional model that considers the keyhole pores formation and trapping. Kan et al.¹² conducted a critical review on the effects of process-induced porosity on the mechanical properties of four alloys commonly fabricated by LPBF: Ti-6Al-4 V, Inconel 718, AISI 316L, and AlSi10Mg. They summarized the existing literature on the tensile properties, fatigue life, impact and fracture toughness, creep response, and wear behavior of these alloys with different levels of porosity. While these models have helped the field understand pore formation, it remains a challenge to translate this knowledge into accurate in-process LPBF monitoring systems.

In-process monitoring of the LPBF process is essential to ensure quality and reliability. One of the most promising approaches is to use high-speed cameras that capture the thermal radiation emitted by the molten pool and the surrounding powder bed.¹³ These cameras can operate in different spectral ranges, such as visible or short-wave infrared (SWIR), and provide rich information about the temperature, shape, size, and dynamics of the molten pool, as well as the spatter, porosity, and surface roughness of the solidified part.¹⁴ However, the raw images from the cameras are often noisy, distorted, and complex. They therefore require advanced feature extraction and machine learning techniques in order to collect meaningful and relevant information that can be used for quality assessment and control.¹⁵ There are also many challenges and open questions that need to be addressed, such as: how to select the optimal spectral range and camera settings; how to deal with noise and distortion in the images; how to choose the best features and algorithms for different applications; how to handle large amounts of data and computational resources; how to validate and interpret the results; and how to integrate the information from multiple cameras and sensors and finally how to leverage extracted sensory features to make fast and accurate predictions of defects, e.g., porosity, while printing an object. These are some of the topics that will be discussed in this research.

In particular our main contributions are:

- We utilize unsupervised variational autoencoders (VAEs) to automatically learn and fuse melt pool features at high speeds up to (20 kHz) from videos recorded by each one of our high-speed camera sensors (i.e., the SWIR and visible wavelength range cameras).
- We demonstrate a physical interpretation for the VAE generated features.
- We train an agile prediction model to utilize features from each sensor independently, as well as their fused feature sets, and predict the occurrence of porosity.
- A prediction performance comparison between these models is provided to highlight the differences of the information provided by each sensor for the pore prediction task.


2. MATERIALS AND METHODS

2.1 Specimen description

A test specimen was 3D printed on a 3D Systems ProX DMP320 laser powder bed fusion (LPBF) machine in 316L stainless steel. This machine employed a 500 W power-adjustable IPG fiber laser to melt metal powder layers with configurable thickness between 30 and 60 μm . Each of these print layers were created by several



(a) printed specimen



(b) Position of the installed cameras

Figure 1: **(a)** An image of the specimen after printing, before detachment from the print bed. **(b)** Top view of the printer with off-axis visual and SWIR cameras installed in the left and right positions respectively to record the melt pool

laser scan lines. For our specimen, layer thickness was set at $30\ \mu\text{m}$, with a laser focal spot of $75\ \mu\text{m}$ and hatching distance $100\ \mu\text{m}$. The powder particle size distribution was between $20\text{-}50\ \mu\text{m}$. The build chamber is vacuum-cycled and filled with an argon over-pressure before commencing the LPBF process.

The specimen (see Figure 1a) consists of a main cylinder (diameter $5\ \text{mm}$ and height $16\ \text{mm}$) with several smaller cylinders protruding at angles of 45 degrees. These smaller cylinders allow easy alignment with CT reconstructions, taking into account possible shrinking. The main cylinder was printed on a squared non-symmetrical base plate that also allows for easy orientation of the object. The base plate itself was printed above $1.8\ \text{mm}$ offsets which were removed mechanically with a cutting disk after completion.

The machine is equipped with the Materialise Control Platform (MCP), which allows us to alter the print parameters for every line. The bulk of the material is printed with nominal settings, i.e., laser power $P = 215\text{W}$ and laser speed $V = 900\text{mm/s}$. Hatching rotates with 67 degrees between layers.

Error layers (i.e. layers printed with off-nominal laser speed and power) are introduced in the object with at least 30 nominal layers (i.e. layers printed with nominal print settings) in between. The interline distance is the hatching distance, $100\ \mu\text{m}$. This spacing between error lines in two dimensions in our experimental setup allows us to identify each introduced porosity with the error line that caused it, with the assumption that nominal settings do not introduce porosity in the bulk of the object. The details about the induced error layers including layer number and laser power and speed can be seen in Table 1. A more detailed description of the specimen is available in Vandecasteele et al.¹⁶

2.2 Printing process monitoring setup

The melt pool was monitored during the printing extending a custom setup explained in Ahar et al.¹⁷ Here we utilized two off-axis high-speed cameras both placed with an inclination angle of about 25° with respect to the normal of the build plate (See Figure 1b). The first camera was an industrial prototype recording in the visible spectrum and its specs can be found in Nourazar et al.¹⁸ The second camera was the Xenics Cheetah 640 TE1 recording in the SWIR range.¹⁹ This off-axis monitoring setup is particularly advantageous for recording the spatter signature in the vertical direction and being in a stationary position independent from the movements in optical path of the laser beam. The video stream is recorded for both cameras at nominal sampling rate of $20,000$ frames per second with each frame having 96×96 pixels and 170×170 pixels for the visual and SWIR cameras respectively. To protect the sensors from the excessive brightness of laser, a short-wavelength pass filter with a cut-off value of $975\ \text{nm}$ was placed before the visual camera, and a high-wavelength pass filter with a cut-off value of $1100\ \text{nm}$ was placed before the SWIR camera. The exposure time of the cameras were set at $30\ \mu\text{s}$. The aperture and shutter time are set manually to optimize the visibility of the spatters while recording the melt pool. Figure 2, demonstrates a sample image frame recorded by each one of the cameras with an 8-bit linear color mapping.

Table 1: Error layers with off-nominal laser settings. The lack-of-fusion and keyhole densities (computed from the CT scan) for each error layer are included as well.

	Starting layer	Laser Power (W) Nominal=215	Laser Speed (m/s) Nominal=900	LoF density	Keyhole density
1	97	215	720	0	0.1913
2	130	215	600	0	0.4747
3	163	215	450	0	2.2515
4	220	215	360	0	3.9637
5	253	215	300	0	4.9502
6	286	215	225	0	5.3807
7	319	322.5	900	0	0
8	376	430	900	0	0
9	409	500	837	0	0.5180
10	442	215	1200	0	0
11	475	215	1800	1.9354	0
12	532	161.25	900	0	0


The video recording was accompanied with the metadata provided by the MCP which was collected at 100 kHz. The metadata included the x, y coordinates of the laser on the build area, the laser on/off signal, the laser power, the laser speed, and timestamps. The high-speed cameras were connected via Camera Link to NI PXI. Data was stored to an NI SSD drive with a high-speed data transfer rate of around 1.5 GB/s.

2.3 Structure of the input data

The proposed methods consist of three machine learning models: two variational autoencoders (VAE) to perform image feature extraction (one VAE per camera), and one regression model for porosity prediction. These models will be trained of the training data described below.

To prepare the training data for the unsupervised VAE learning of image features, we synchronize and annotate each one of the captured images of both cameras with the corresponding metadata collected from the MCP controller. Then, we randomize their capture order and two separate subsets of images for training and inference are selected from the total image pools. The training subset includes a total of 622,592 images. To further increase variability of the training set, print recording from two additional printed objects were also added to the pool. Though since those two objects were not CT scanned, we did not include their data for training and analyzing the final porosity prediction model. During training and inference, the images undergo preprocessing. The unprocessed images are cropped to a 40x40 pixel region centered around the melt pool's center.¹³ Cropping and centering aim to eliminate extraneous data and speed up processing.

In order to create training and test sets for the final porosity prediction model, all extracted image features are associated to the extracted pore densities from the CT scan data. Here we create our model dataset by choosing a subset of features selected only from the 12 error layers (print layers with imposed off-nominal settings) and then add features from 3 layers with nominal laser settings. This way we created a better balance in the quantity of defective and non-defective values, with an emphasis to be more sensitive on detecting defective print conditions. Our initial experiments utilizing data from all print layers which included many more print layers with nominal settings (+500 layers) revealed that model loss-function optimization would be biased toward better predicting target values for non-defective data points and consequently reduced sensitivity for predicting target values for porosity. The order of the training data was then randomized. For each experiment, a subset of 70% of the training samples were used to train the porosity prediction model, while the remaining 30% comprised the validation set. The test set then comprised of the remaining data from all print layers, including those layers not ever shown to the model during training phase. To aid model training, all features were centered by dividing to their average values.



(a) Visual camera image

(b) SWIR camera image

Figure 2: (a) A sample frame captured with visible spectrum range camera, and (b) A sample frame captured with SWIR camera. The SWIR image has been cropped and resampled to show a similar field of view as the visual camera image.

2.4 Image Feature Extraction and Fusion

Leveraging a subset of the data, we train two fully-connected VAEs to extract meaningful features from both input modalities. As depicted in Fig 3, the neural network architecture comprises a probabilistic encoder, $e_\phi(\mathbf{x})$, which compresses the input image \mathbf{x} into a lower-dimensional probabilistic representation \mathbf{Z} (feature space). This representation is modeled as a multivariate Gaussian probability distribution, characterized by a mean vector $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$: $\mathbf{Z} = e_\phi(\mathbf{x}) = \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$. Subsequently, the decoder $d_\theta(\mathbf{z})$ reconstructs the output image $\hat{\mathbf{x}}$ from the features, \mathbf{z} sampled from the multivariate Gaussian distribution \mathbf{Z} .

Each layer of the VAE employs a fully-connected architecture, with the exception of the connections between the multivariate Gaussian, with its mean vector $\boldsymbol{\mu}$ and variance vector $\boldsymbol{\sigma}$, and a feature vector \mathbf{z} . This connection utilizes a sampling procedure. Following each fully-connected layer, batch normalization and a rectified linear activation function (ReLU) are applied, except for the encoder's final layer. The layer dimensions are specified at the top of Fig. 3. The VAE is optimized using a loss function (Eq. 1) that balances reconstruction accuracy and feature space smoothness:

$$L(\mathbf{x}, \hat{\mathbf{x}}) = \|\mathbf{x} - d_\theta(e_\phi(\mathbf{x}))\|_2^2 + \beta KL(e_\phi(\mathbf{x}) \|\mathcal{N}(\mathbf{0}, \mathbf{I})) \quad (1)$$

where $KL(\cdot)$ denotes the KL-divergence loss which is used to enforce smoothness in the feature space. Inspired by the β -VAE,²⁰ we introduce an extra scalar β that dynamically adjusts the KL-divergence loss using a ControlVAE²¹ framework that steers the KL-divergence loss towards a specific set value using a non-linear PI-controller. This approach has been shown to improve reconstruction performance while maintaining comparable KL-divergence values.

The VAE models were optimized using the Adam optimizer²² with a learning rate of 0.0005, a batch size of 64, and a training duration of 30 epochs. Following each epoch, the learning rate was dynamically adjusted using an exponential decay factor of $\gamma = 0.99$. The KL-divergence loss target was set to 40, and the PID controller was updated after each epoch.

During training, the target output is set to match the input, encouraging the network to learn a probabilistic lower-dimensional representation \mathbf{z} that effectively summarizes the input image. The mean vector of the lower-dimensional representation then serves as a set of features that capture the essential characteristics of the input datum. The incorporation of a probabilistic encoder and the KL-divergence term distinguishes a VAE from a regular autoencoder (AE). The probabilistic encoding promotes smooth variations in reconstructed images. Simultaneously, the KL-divergence regularizer drives the encodings to organize themselves in the feature space, leading to several benefits:

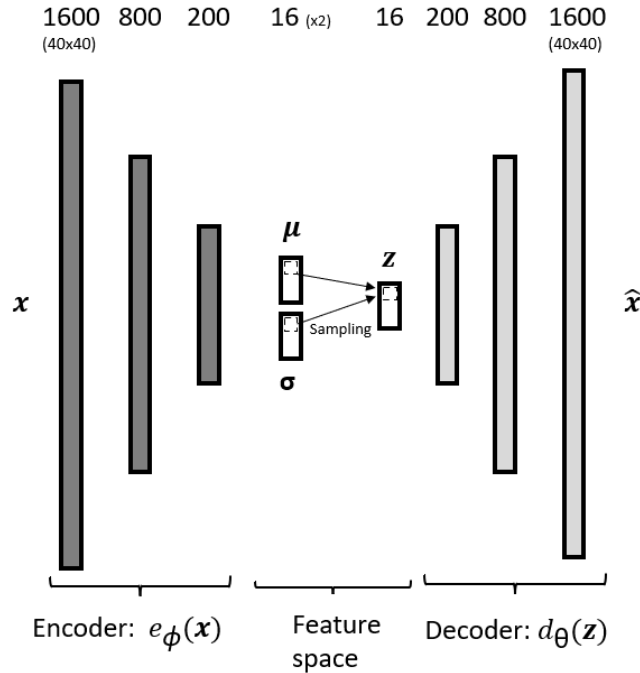


Figure 3: The architecture of the VAE.

- Enhanced Semantic Interpretation: Neighboring encodings correspond closely in the input space, facilitating meaningful interpretation of the acquired features.
- Improved Disentanglement: Each dimension in the feature space is strongly correlated with a single factor of variation in the input space, enabling disentanglement of the underlying factors of variation.
- Enhanced Interpolation Capabilities: The smooth transitions in the feature space allow for effective interpolation between data points.

These advantages collectively contribute to greater interpretability of the acquired features, facilitating comprehensive analysis and furthering our understanding of the underlying data.

During defect detection, only the encoder parts of the VAEs are used to process the image. At that point the encoders serve merely as feature extractors. The corresponding two sets of features, one for each modality, are then input into the defect detection model. The combination, or fusion, of these features is a simple concatenation.

Leveraging the inherent smoothness of the VAE's feature space, we can randomly sample from the feature space, decode the samples, and generate corresponding samples in the image space. This enables us to analyze the role of each feature by systematically varying its value while keeping other features constant, and then observing the resulting image-space variations. Preferably, the sequence of images generated per feature should exhibit a meaningful relationship with physical properties of the melt pool, facilitating interpretation.

To validate this concept, we generated and decoded a sequence of samples for each feature in the feature space. For each sample in the sequence corresponding to a feature dimension, we constructed a vector where each dimension value is set to 0 except for the element of that dimension, which was set to a specific value. This was ranged from -3 to 3 standard deviations in increments of 0.5 . The corresponding decoded images were then visualized. This was done for both imaging modalities.

2.5 Prediction model

In these experiments, complex models and deep neural networks are categorically ignored mainly due to strict run-time limitations of the problem.¹⁸ Instead we utilize a compact model which was previously demonstrated

to make reliable predictions.¹³ The model is a 2 layer perceptron consisting of 2 fully connected dense layers with 164 neurons in the hidden layer, a sigmoid activation function in the hidden layer, and a rectified linear unit (ReLU) activation function in the output. The model is implemented in python language using Keras neural network library.²³ To facilitate better comparison conditions, the model is trained with 100 epochs with the Adam optimizer and the mean absolute error loss function. These settings remain constant for separate training sessions conducted for each camera sensor and their combinations, which hereafter we refer to them as VIS, SWIR, and SWIRVIS for those models trained with features from the visible wavelength range camera, SWIR camera, and both together respectively.

3. EXPERIMENTAL RESULTS

In this section first we explain the results of post-print scanning of the test object and analysis of the created pores in response to the tested laser settings. This analysis effectively generates the ground-truth data for evaluation of prediction model. Next, extracted VAE features are demonstrated and analyzed. Finally, porosity prediction results are then presented and analyzed.

3.1 CT segmentation

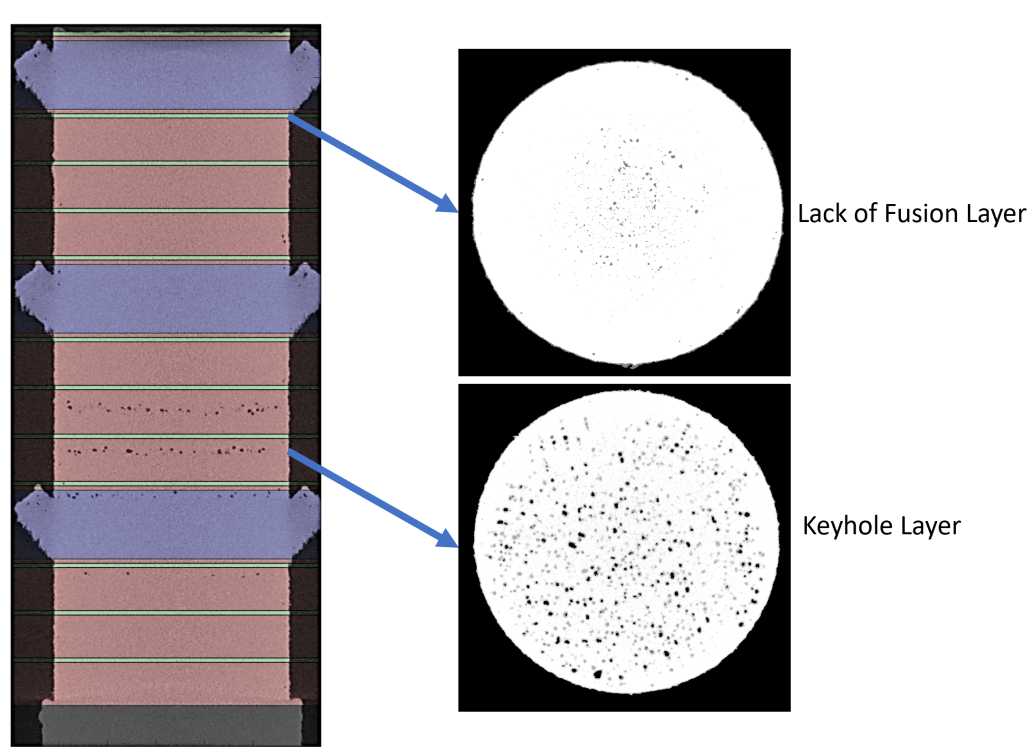



Figure 4: Left side: Illustration of Vertical slice of the cylinder CT scan aligned to the CAD model. The overlaid green horizontal lines represent the position of print layers with off-nominal laser power and speed, while red and blue regions represent those print layers with nominal setting in the cylinder bulk. Top right: A sample horizontal slice of layer 410 printed with 2x nominal laser speed which created the lack of fusion. Bottom right: A sample horizontal slice below layer 186 printed with 1/3x nominal laser speed which created keyhole porosity.

After printing, detachment from the printer bed, and grinding and polishing to remove the supports, the 3D printed cylindrical object was CT scanned resulting in a CT reconstruction with isotropic 10 μm voxel resolution. The CT scan was denoised with a total variation denoising method²⁴ and thresholded to create a pore segmentation.



(a) VIS features only (b) SWIR features only (c) SWIR and VIS features together

Figure 5: Final porosity prediction results averaged over each print layer in our test set. Ground truth pore density values for each error layer are indicated with circle markers as well. (a) Prediction results based on model trained with only VIS features. (b) Prediction results based on model trained with only SWIR features. (c) Prediction results based on model trained with both SWIR and VIS sensor fused features.

Registration of the CT reconstruction with the experiment’s coordinate system was performed by minimizing the deviation of several features from their theoretical values with respect to translation and rotation degrees of freedom. A side view of the aligned CT scan with the CAD model is shown in Figure 4. The error lines that created porosity are indicated as overlaid green lines. When keyhole pores are created, significant differences can be observed on the number and depth of the pores. This is to be expected: the penetration depth of the laser and the subsequent keyhole pore creation processes show a complex dependence on the laser parameters.²⁵

Table 1, summarizes main characteristics of the error layers and presents the extracted target pore densities for the lack of fusion and keyhole porosity. Note that not all the off nominal laser settings resulted in inducing porosity on the printed object. This is particularly important for the porosity prediction task because the prediction model has to learn the real transitions from conduction mode to the pore creation modes instead of learning to merely react on any detection of variation in laser parameters and corresponding melt pool characteristics.

3.2 Porosity prediction results

After separately training the prediction model with VIS, SWIR and SWIRVIS features, the trained models are tested over the test set mentioned in section 2.3. Figure 5 shows the test results for each set of sensor features averaged per layer. This way, the average of the predictions is depicted per layer for the density of the LoF or Keyhole pores, densities whose ground truth values were extracted from the CT scan analysis.

It is immediately clear from the VIS-only prediction results that while the model been able to accurately predict the density of both LoF (shown in red color plot) and keyhole (shown in green color) for the layers of which they are generated, the keyhole predictions in the non-defective layers (essentially containing no pore) appear to be highly noisy including lots of false positive cases. Initially, we expected that this over sensitivity issue of the VIS model was due to sub-optimal performance of the VAE in generating features which encode the melt pool information. However, our multiple attempts to regenerate VIS features (not reported in this paper) while fine tuning the VAE parameters, revealed that this issue is mainly inherent to the quality of the image data and potentially can be improved by better camera sensor calibration. Nonetheless, we did not further investigate hardware related solutions which will remain for future work.

Figure 5b shows the prediction result for the model trained with the SWIR features. Two important differences can be observed when compared with the VIS-only case of Figure 5a: First, the SWIR model fails to react to the lack of fusion layer while more accurately predicting the pore densities in layers generating keyhole pores. This demonstrates that rise of the laser energy density when transitioning from the conduction mode to the keyhole generation mode is accurately captured via its thermal signatures (e.g., visible increase in size of the melt pool thermal emission, spatters, cool down rate of the melt pool trail, etc.), while the drop in laser thermal footprints when transitioning to the LoF generation mode seems not to be as evident, and thus not being picked

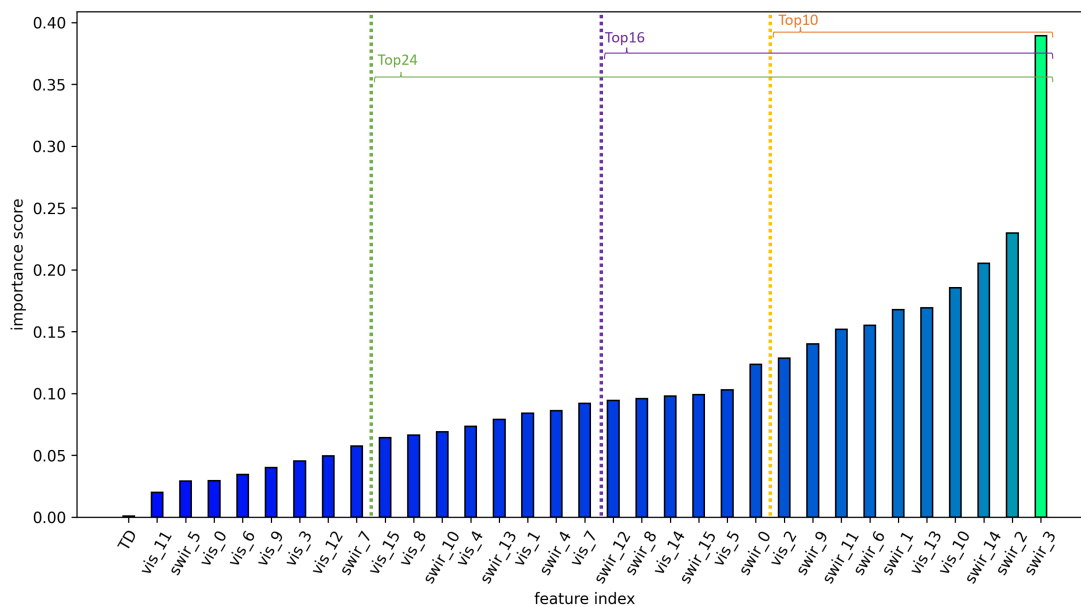


Figure 6: Importance Analysis of the VAE features based on their contribution to the final porosity predictions

up by the VAE encoded features. Second, compared to the VIS-only predictions, keyhole density predictions in non-defective layers are significantly less noisy and more stable. This again demonstrates the clear distinction the prediction model can make between thermal signatures of melt pool when transitioning to the keyhole regime. Further investigation on the VAE features is required to verify our reasoning, investigations which we plan to pursue later.

Arguably the best prediction results can be observed when the model is trained on fusion of both SWIR and VIS features. Figure 5c, demonstrates the SWIRVIS model results. It is interesting to see that the prediction model has been able to capture the positive aspects of each feature set and provide not only accurate predictions for both lack of fusion and keyhole layers but also correctly identify the non-defective layers with very low noise.

3.3 Analysis of the results

After successful demonstration of prediction results from sensor-fused VAE features, we would like to further investigate the individual contribution of each feature to the final prediction performance. To do so, we conduct a permutation importance analysis.²⁶ Permutation importance is a model inspection technique that can be used for any fitted estimator when the data is tabular. It is a method that shuffles the values of each feature and measures its impact on the model's final performance. The more the performance drops, the more important the feature is. Permutation importance is model-agnostic and can be calculated many times with different permutations of each feature. For our case, we conduct the analysis with 10 different permutations for each SWIR and VIS feature on the SWIRVIS trained model over the test set. Figure 6 shows the final calculated importance scores averaged over the 10 trials. The standard deviation of the calculated scores for all features is 0.04 ± 0.02 which represents good reliability of the assigned scores. Features in Figure 6 are sorted and color coded based on their mean importance scores for improved readability. Looking at the top 10 important features, it is clear that SWIR sensor features are dominant (7 out of 10) and have more significant role in accurate porosity predictions overall, with feature "SWIR_3" by a large margin being chosen as the most important feature, followed by "SWIR_2" and "SWIR_14".

Fig 7 shows some example images that were sampled from the top five most important VAE features using the earlier described method. From top to bottom, these represent the 3rd SWIR feature, 2nd SWIR feature, 14th SWIR feature, 10th VIS feature, 13th VIS feature. Seemingly, the 3rd SWIR feature is highly correlated to




Figure 7: Resulting sequence of images (horizontally) of 5 features after decoding specifically created feature samples. Top to bottom: 3rd SWIR feature, 2nd SWIR feature, 14th SWIR feature, 10th VIS feature, 13th VIS feature.

the size of the melt pool. It is thus not unsurprising that "SWIR_3" is significant in predicting porosity. Melt pool size is in turn heavily correlated to energy density and thus keyhole, nominal and LoF printing modes.⁹

Meanwhile, the 2nd SWIR also seems to be somewhat related to the size as well as to the size and direction of the melt pool wake. The 14th SWIR feature shares some similarities to the 3th SWIR feature although it appears that the melt pool size remains constant and the streaks appear to be longer and more straight. This could be more correlated to large vapor plumes, mostly associated with high energy densities. Similarly, large melt pool wakes are associated with large energy densities (keyhole), while the lack thereof could indicate low energy densities (LoF).⁹

Finally, the 10th VIS feature looks to be both correlated to the size of the melt pool as well as the occurrence of spatter particles. Similarly, the 13th VIS feature is also correlated to spatter particles, but in addition, to the size and direction of the spatter particles. These results also make sense as the occurrence, size, and direction of spatter particles have been shown²⁷ to be strong indicators for melt pool stability.

These results show that the VAE is capable of producing features that are correlated to the physical phenomena that we expect to occur during printing. However, it is necessary to note that this is not the case for every produced feature. We noticed that some of the less important features appear to be more abstract or noisy, whilst other appear to show no changes at all in the image space. These results may indicate that a more compact representation exists. Although the features show some disentanglement, it is clear that there is still some semantic overlap between features. This can be improved by putting more weight on the KL-divergence, but it should be noted that this weighting may lead to a decreased reconstruction quality since these two properties are in a direct trade-off.

In addition to this feature analysis, we investigated the possibility of reducing the number of fused features based on the importance analysis results. We aim for training our pore prediction model while dropping 25%(8 features), 50% (16 features) and 75% (24 features) of least important features from the model input. However, our initial experiments showed that dropping 75% least important features already causes the model to completely miss on picking up the lack of fusion layer. This was due to the fact that not enough VIS features were kept (only 2 out of 8). We were able to resolve this issue by keeping two more features (the top 10 features in total) which included one more VIS feature as well. Figure 8 demonstrates the prediction performance of these redacted feature sets tested on the test set. Comparing the results with those in Figure 5c as our reference, it is clear that even by dropping 60% of the least important features in case of Figure 8a, we can achieve a close performance both in detecting LoF and keyhole layers as well as stable predictions in the case of non-defective layers. This

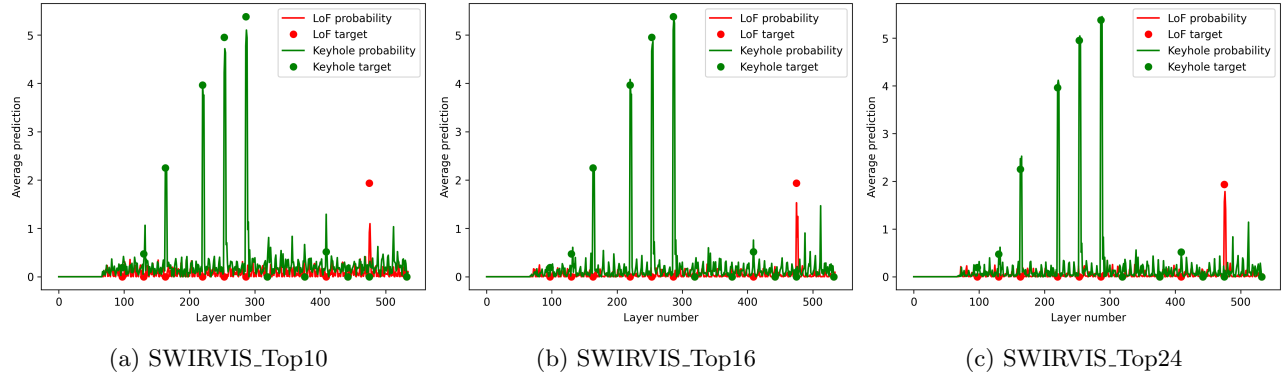


Figure 8: Final porosity prediction results averaged over each print layer in the test set. Ground truth pore density values for each error layer are indicated with circle markers. (a) Prediction results based on model trained with the top 10 important VIS and SWIR features. (b) Prediction results based on model trained with the top 16 important VIS and SWIR features. (c) Prediction results based on model trained with the top 24 important VIS and SWIR features.

accuracy continues to improve by keeping 50% of the most important features in case of Figure 8b to the point that for the case of keeping top 75% most important features shown in Figure 8c, we observe that dropping those 25% least important features actually makes some improvements over the case of using all the features. For example, a false positive Lof detection in layer 75 of Figure 5c no longer shows up in Figure 8c. The keyhole pore density predictions are also equally accurate and stable.

To further quantify our qualitative analysis, we calculate the Pearson Correlation Coefficient (PCC), Spearman Rank order Correlation (SROC) and the Mean Squared Error (MSE) between the set of all predicted pore density values and the target values over the validation set. We also present a layer-wise comparison over the test set where we averaged prediction results per print layer separately for each porosity, and then averaged across porosity types. We also calculated the MSE and normalized it over the maximum ground truth values such that resulted prediction error values are comparable across differently trained models. These results are summarized in Table 2 We observe that all of these quality measures verify our qualitative analysis as well by giving highest correlation scores and lowest prediction errors to the case of SWIRVIS with all 33 features or with selected top 24 important features.

Table 2: Porosity prediction performance of the tested model using different combinations of the VAE features. Pairwise comparisons are performed for all single prediction values over the testset-1. Layer wise comparisons are performed over average of single predictions per print layer in testset-2.

		SWIR_all16	VIS_all16	SWIRVIS_all33	SWIRVIS_top24	SWIRVIS_top16	SWIRVIS_top10
Pairwise comparison	PCC	0.9741	0.9457	0.9839	0.9827	0.9729	0.9493
	SROC	0.8916	0.8862	0.9173	0.9095	0.8981	0.8787
	MSE	0.2204	0.4394	0.135	0.1462	0.2253	0.411
Layerwise comparison	NMSEAll	0.0045	0.06	0.0021	0.0024	0.004	0.008
	NMSElof	0.0124	0.003	0.0025	0.0022	0.0041	0.0108
	NMSEkeyhole	0.0045	0.1189	0.0033	0.004	0.0064	0.0121

4. CONCLUSIONS

In this work, we propose a low-latency LPBF defect detection algorithm based on fusion of images from high-speed cameras in the visible and short-wave infrared (SWIR) spectrum ranges. First, we designed an experiment to print an object while imposing lack of fusion and keyhole porosity defects on certain print layers by changing the laser power and speed parameters. During the print, the laser's melt pool was recorded with the two high-speed cameras at nominal sampling rates of 20 kHz. We utilized a variational autoencoder to automatically extract two sets of meaningful features from the corresponding raw sensor images. The melt pool image features are then annotated with each layer pore densities extracted from the CT scan of the printed object. These

annotations are further utilized as ground truth for supervised training and evaluation of a compact neural network model to predict the occurrence of porosity from the fused features. We compared the prediction performance of our sensor fused model with models trained on image features from each camera separately. We observe that the SWIR imaging is sensitive to keyhole porosity while the visible-range optical camera better characterizes lack-of-fusion porosity. By fusing features from both cameras, we are able to accurately predict both pore types, thus outperforming both single camera systems. We also conducted a feature importance analysis and demonstrate that by redacting 25% of the fused features based on their importance scores, we can achieve comparable performance thus further reduce the computational load of the proposed monitoring pipeline. It should be noted that after completing the training phase both for the unsupervised training of the VAE model as feature extractor, and the supervised training of the prediction model, our proposed setup can achieve in-situ monitoring with a similar amount of computations as an existing real-time monitoring system,¹⁸ suggesting that the defect prediction latency of this monitoring system is comparable to the state-of-the-art. This brings us one step closer to create a closed-loop control to automatically steer the laser settings for optimal defect-free LPBF printing.

ACKNOWLEDGMENTS

This research is supported by the Flemish Innovation and Entrepreneurship Agency through the imec ICON research project MultiPLICITY (HBC.2022.0094) and by Flanders Make (the strategic research center for the manufacturing industry) & imec (the international research & development organization, active in the fields of nanoelectronics and digital technologies). This research received partial funding from the Flemish Government under the “Onderzoeksprogramma Artificiele Intelligentie (AI) Vlaanderen” programme. We would like to express our gratitude to other partners of the project consortium for their assistance in technical discussions including Flanders MAKE@KU Leuven; Materialise; imec Vision Lab, University of Antwerp; Dekimo; Manufacturing Processes and Systems (MaPS)@KU Leuven; Xenics; Rejig; and AdditiveLab.

REFERENCES

- [1] Ian Gibson, I. G., “Additive manufacturing technologies 3d printing, rapid prototyping, and direct digital manufacturing,” (2015).
- [2] Wong, K. V. and Hernandez, A., “A review of additive manufacturing,” *International scholarly research notices* **2012** (2012).
- [3] Gu, D. D., Meiners, W., Wissenbach, K., and Poprawe, R., “Laser additive manufacturing of metallic components: materials, processes and mechanisms,” *International materials reviews* **57**(3), 133–164 (2012).
- [4] Kruth, J.-P., Levy, G., Klocke, F., and Childs, T., “Consolidation phenomena in laser and powder-bed based layered manufacturing,” *CIRP annals* **56**(2), 730–759 (2007).
- [5] Murr, L. E., Gaytan, S. M., Ramirez, D. A., Martinez, E., Hernandez, J., Amato, K. N., Shindo, P. W., Medina, F. R., and Wicker, R. B., “Metal fabrication by additive manufacturing using laser and electron beam melting technologies,” *Journal of Materials Science & Technology* **28**(1), 1–14 (2012).
- [6] Gong, H., Rafi, K., Gu, H., Starr, T., and Stucker, B., “Analysis of defect generation in ti-6al-4v parts made using powder bed fusion additive manufacturing processes,” *Additive Manufacturing* **1**, 87–98 (2014).
- [7] Cherry, J., Davies, H., Mehmood, S., Lavery, N., Brown, S., and Sienz, J., “Investigation into the effect of process parameters on microstructural and physical properties of 316l stainless steel parts by selective laser melting,” *The International Journal of Advanced Manufacturing Technology* **76**, 869–879 (2015).
- [8] Nudelis, N. and Mayr, P., “A novel classification method for pores in laser powder bed fusion,” *Metals* **11**(12), 1912 (2021).
- [9] Thanki, A., Jordan, C., Booth, B., Verhees, D., Heylen, R., Mir, M., Bey-Temsamani, A., Philips, W., Witvrouw, A., and Haitjema, H., “Off-axis high-speed camera-based real-time monitoring and simulation study for laser powder bed fusion of 316l stainless steel,” *The International Journal of Advanced Manufacturing Technology* **125**, 4909–4924 (2023).
- [10] du Plessis, A., “Effects of process parameters on porosity in laser powder bed fusion revealed by x-ray tomography,” *Additive Manufacturing* **30**, 100871 (2019).

- [11] Wang, W., Ning, J., and Liang, S. Y., “Analytical prediction of keyhole porosity in laser powder bed fusion,” *The International Journal of Advanced Manufacturing Technology* **119**(11-12), 6995–7002 (2022).
- [12] Kan, W. H., Chiu, L. N. S., Lim, C. V. S., Zhu, Y., Tian, Y., Jiang, D., and Huang, A., “A critical review on the effects of process-induced porosity on the mechanical properties of alloys fabricated by laser powder bed fusion,” *Journal of Materials Science* **57**(21), 9818–9865 (2022).
- [13] Booth, B. G., Heylen, R., Nourazar, M., Verhees, D., Philips, W., and Bey-Temsamani, A., “Encoding stability into laser powder bed fusion monitoring using temporal features and pore density modelling,” *Sensors* **22**(10), 3740 (2022).
- [14] Khan, A., Jaffery, S. H. I., Hussain, S. Z., and Dilawar, S., “Numerical and experimental characterisation of melt pool in laser powder bed fusion of ss316l,” *Research Square preprint* (2023).
- [15] Wang, C., Tan, X., Tor, S., and Lim, C., “Machine learning in additive manufacturing: State-of-the-art and perspectives,” *Additive Manufacturing* **36**, 101538 (2020).
- [16] Vandecasteele, M., Heylen, R., Iuso, D., Thanki, A., Philips, W., Witvrouw, A., Verhees, D., and Booth, B. G., “Towards material and process agnostic features for the classification of pore types in metal additive manufacturing,” *Materials & Design* **227**, 111757 (Mar. 2023).
- [17] Ahar, A., Heylen, R., Verhees, D., Blanc, C., and Bey-Temsamani, A., “Accelerated monitoring of powder bed fusion additive manufacturing via high-throughput imaging and low-latency machine learning,” in [*IFIP International Conference on Artificial Intelligence Applications and Innovations*], 250–265, Springer (2023).
- [18] Nourazar, M., Booth, B., and Goossens, B., “A GPU optimization workflow for real-time execution of ultra-high frame rate computer vision applications,” *Journal of Real Time Image Processing* **21**(5), 13 (2024).
- [19] “Short-wave infrared imagers — Xenics.” <https://www.xenics.com/short-wave-infrared-imagers/cheetah-series/> (Accessed: 2023-12-01).
- [20] Higgins, I., Matthey, L., Pal, A., Burgess, C., Glorot, X., Botvinick, M., Mohamed, S., and Lerchner, A., “beta-VAE: Learning Basic Visual Concepts with a Constrained Variational Framework,” (Nov. 2016).
- [21] Shao, H., Yao, S., Sun, D., Zhang, A., Liu, S., Liu, D., Wang, J., and Abdelzaher, T., “ControlVAE: Controllable variational autoencoder,” in [*Proceedings of the 37th International Conference on Machine Learning (ICML)*], (2020).
- [22] Kingma, D. P. and Ba, J., “Adam: A Method for Stochastic Optimization,” (Jan. 2017).
- [23] Chollet, F. et al., “Keras,” (2015). Available at <https://github.com/fchollet/keras>.
- [24] Heylen, R., Thanki, A., Verhees, D., Iuso, D., Beenhouwer, J. D., Sijbers, J., Witvrouw, A., Haitjema, H., and Bey-Temsamani, A., “3d total variation denoising in x-CT imaging applied to pore extraction in additively manufactured parts,” *Measurement Science and Technology* **33**, 045602 (jan 2022).
- [25] Zhao, C., Parab, N. D., Li, X., Fezzaa, K., Tan, W., Rollett, A. D., and Sun, T., “Critical instability at moving keyhole tip generates porosity in laser melting,” *Science* **370**(6520), 1080–1086 (2020).
- [26] Breiman, L., “Random forests,” *Machine Learning* **45**, 5–32 (2001).
- [27] Heeling, T., Gerstgrasser, M., and Wegener, K., “Investigation of selective laser melting spatter characteristics for single- and multi-beam strategies using high speed imaging,” (2017).