

IN-SITU ANOMALY DETECTION IN ADDITIVE MANUFACTURING WITH GRAPH NEURAL NETWORKS

Sebastian Larsen & Paul A. Hooper

Department of Mechanical Engineering

Imperial College London

London, SW7 2BX, UK

{s112215, paulhooper}@imperial.ac.uk

ABSTRACT

Transforming a design into a high-quality product is a challenge in metal additive manufacturing due to rare events which can cause defects to form. Detecting these events in-situ could, however, reduce inspection costs, enable corrective action, and is the first step towards a future of tailored material properties. In this study a model is trained on laser input information to predict nominal laser melting conditions. An anomaly score is then calculated by taking the difference between the predictions and new observations. The model is evaluated on a dataset with known defects achieving an F1 score of 0.821. This study shows that anomaly detection methods are an important tool in developing robust defect detection methods.

1 INTRODUCTION

Machine learning (ML) methods are increasingly used in additive manufacturing (AM) to automate inspection (Scime & Beuth, 2018), predict defects (Larsen & Hooper, 2022), as well as potentially tailor microstructure and material properties (DebRoy et al., 2021). In Laser-Powder Bed Fusion (L-PBF), a thin layer of powdered material is spread across a build plate. A laser then scans back and forth melting the powder. The process is repeated layer-by-layer until a 3-dimensional part is formed. As a result of chaotic melting conditions, micron sized flaws can occur. Hence, in-situ monitoring cameras have been combined with ML to observe the process. There is, however, a high cost associated with collecting labelled data for training ML algorithms. This requires X-ray computed tomography (XCT) to characterise the flaws which form as pores inside the material. XCT is also limited to very small parts. It is, therefore, desirable to develop robust anomaly detection approaches which require minimal flaws for training ML models.

Another challenge is to represent the AM process complexity with the available information. This can be addressed by geometric deep learning methods. Graph neural networks (GNN) have become increasingly popular for modelling data that takes on a graph type structure. This includes molecules (Gainza et al., 2020), physics simulations (Sanchez-Gonzalez et al., 2020), traffic prediction (Derrow-Pinion et al., 2021), and material property prediction (Xie & Grossman, 2018). Monitoring signals are dependent on the geometry of the component being manufactured as well as the laser scanning strategy. Therefore, a graph structure is highly suited to modelling the underlying process physics.

The objective of this study is to develop an anomaly detection approach that requires no labelled data other than nominal processing conditions. Due to spatio-temporal correlations in the data, a graph structure is selected to represent the positional information. Given a graph and the laser feature inputs, such as power and scan direction, the model is trained to predict nominal melting conditions. This serves as an approach to compare new observations to the expected result, yielding an anomaly score.




Figure 1: (a) Representative image of the melt pool which forms as the laser melts material. (b) Segmented image. (c) Temporal signal of a melt pool and laser power. (d) Laser scan direction feature. (e) X-ray CT scan image slice with pore defects shown as dark spots.

2 METHODOLOGY

2.1 SENSOR AND DATA COLLECTION

Images were acquired of the melt pool with a high-speed recording at 100 kHz (Fig. 1a) (Hooper, 2018). The camera is located co-axial to the laser and measures the melt pool radiation. The images were segmented to extract a number of melt pool features including: size, shape, intensity, and number of spatter particles (Fig. 1b). These features are used as node labels described in the next section. The melt pool signal correlates with the laser power which is pulse width modulated (Fig. 1c). The laser also changes direction as it scans over the layer (Fig. 1d). A number of defects were injected into a part during processing and the measured signal taken from the next layer. After the component was built, an X-ray CT scan was carried out (Fig. 1e). Defects, given by porosity, were characterised within the part using a segmentation algorithm. A small amount of anomalous process conditions could then be mapped back to the measured signal location to act as an evaluation dataset for the anomaly detector.

2.2 MACHINE LEARNING

The framework for anomaly detection can be seen in Fig. 2. An anomalous observation is defined as a data point that differs considerably from the other data within a set of observations. For instance reduced intensity or excessive spatter shown in Fig. 1. Therefore, the goal is to learn a score function that can be compared to new observations. Since there are contextual events, such as the laser features and part geometry, this known physics should be incorporated into the model. Hence, we propose to predict the expected sensor observation (\hat{Y}) given a set of laser inputs (X), $\hat{Y} = f_{\theta}(X, A)$, where f_{θ} is a function mapping of inputs to outputs parametrised by θ and A is the adjacency matrix. At test time the predicted features are compared to the measured features. This is computed with an error metric given by

$$Z_i = \sum_{j=1}^n (\hat{Y}_{i,j} - Y_{i,j}), \quad (1)$$

where Z will approach a Gaussian distribution for nominal conditions compared to anomalies. Laser features including scan direction, power, node number and track number were used as model inputs (node features).

As the laser data is correlated both temporally and spatially, a graph was constructed based on the nearest neighbours of each node. These are expected to correlate since a defect can span multiple laser tracks. The input information is fed to a Graph Transformer network (Shi et al., 2020). This enables the use of edge attributes as well as node features to be represented. In this case, we label edges based on whether they are connected to an adjacent scan track or a not, i.e. temporally or spatially correlated. This is followed by a Fully Connected (FC) post-message passing layer before predicting the expected sensor signature.

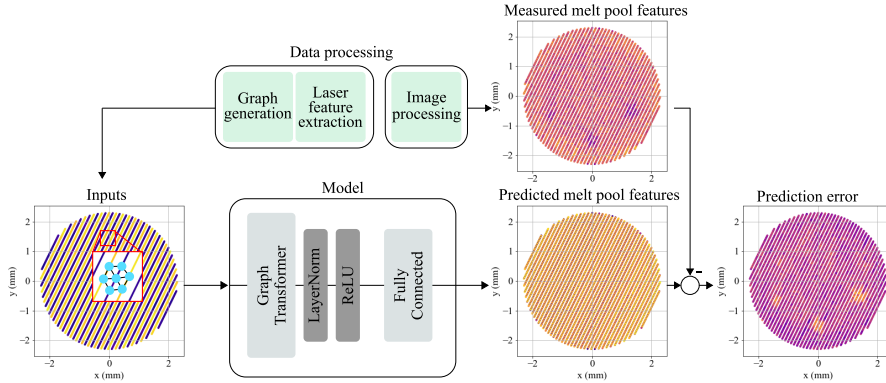


Figure 2: The framework for anomaly detection developed in this study.

The model was trained on nominal processing conditions to minimise the mean squared error (MSE) using the Adam optimiser (Kingma & Ba, 2014). Each graph corresponded to a layer of deposited material, with training data on eight graphs totalling 182,450 nodes. The model was evaluated on four graphs known to contain anomalies, where 1,735 nodes were defective and 70,495 were nominal with a data imbalance of 40.6 to 1.

After the error score was computed it was smoothed to remove noise and to enable information propagation in the local neighbourhood of each node. This was performed with a symmetric normalisation and aggregation given by

$$Z' = D^{-\frac{1}{2}} A D^{-\frac{1}{2}} Z, \quad (2)$$

where D is the degree matrix.

3 RESULTS AND DISCUSSION

3.1 NODE LEVEL PREDICTIONS

A threshold was obtained for the outlier model using the precision-recall curve which was chosen due to the data imbalance (Fig. 3a). The area under the curve, or average precision (AP) was 0.832, while a maximum F1 score was measured as 0.821. Fig. 3b shows a histogram of nominal and anomalous prediction errors along with the selected threshold from the precision-recall curve.

The prediction error was expected to converge to a Gaussian distribution if the sources of variation are accounted for. This motivates studying the selected threshold from a quantiles perspective (Fig. 3c). The prediction errors were mapped to a standard Gaussian distribution along with the threshold. This enabled an estimate of how many false positives (FP) might be expected for a given nominal layer. In this case, a threshold of 3.42 resulted in a false positive rate (FPR) of around 6 in 10,000 (Table 1). When compared on the quantile plot of a Gaussian distribution it is expected to be closer to 3 in 10,000. Therefore, an improvement may still be possible though this is in good agreement. However, the quantiles of the unseen data are different, particularly in the negative region. This is likely due to the few graphs used in the dataset (only 8 in training). We believe this can be further improved by including more examples thereby improving the prediction model.

The input feature importance was also evaluated (Fig. 3d). The laser power was the most important in making predictions. This is due to the laser turning on and off which correlates with the signal in the camera. When the laser turns off, the signal reduces. This information is essential to include since anomalous events, such as a reduced signal, can also be associated with defect formation. The node and track features were less important but provide some contextual information as the laser traverses the layer heating the component over time. The scanning directions were close to redundant in the predictions.

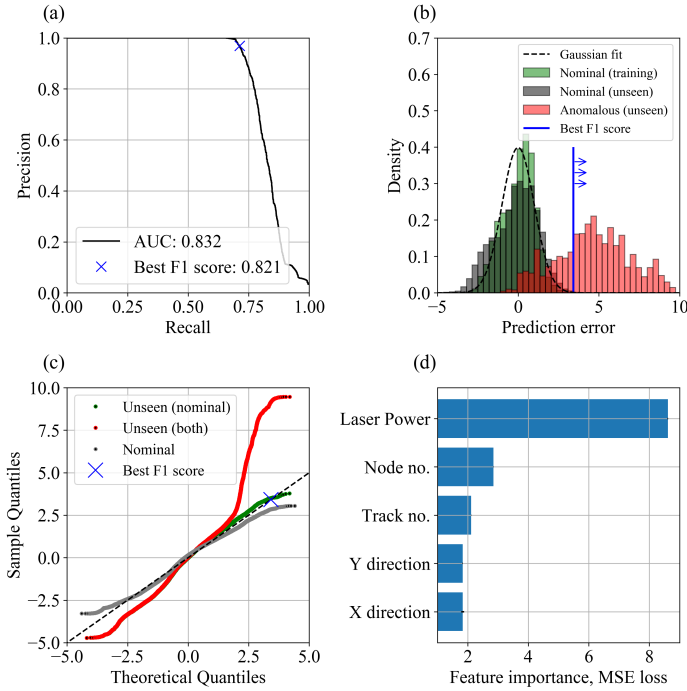


Figure 3: (a) Precision-recall curve. (b) Histogram of error for nominal and anomalous data. (c) QQ plot of both nominal and anomalous datasets corresponding to (b). (d) Feature importance of inputs.

3.2 MODEL COMPARISONS

The model was compared to: Graph Attention Networks (GAT) (Veličković et al., 2017), a Graph Convolutional Network (GCN) (Kipf & Welling, 2016), a Graph Isomorphism Network (GIN) (Xu et al., 2018), and a FC network (Table 1). Furthermore, an autoencoder (AE) anomaly detection method was included and trained to reconstruct both laser and image features.

The Graph Transformer (Graph-T) model had the highest average precision (AP) and F1 score. The closest performing model (GAT) had similar AUROC but lower AP. We consider the AP a better evaluation in this application due to the data imbalance and add larger weight to the lower FP achieved by the Graph-T model. When comparing the FC with the GNN models (GAT, GCN, GIN), the loss was higher for the GNNs. It is thought that these models might over-smooth the inputs causing some information to be lost in the forward pass, while the FC does not utilise the graph structure until test time. The Graph Transformer can, however, take advantage of both pieces of information.

Table 1: Comparative performance of different models including Graph Transformer with an absolute anomaly metric (Graph-T-A) and without output smoothing (Graph-T-Z).

Model	AP	AUROC	F1	FP	FN	TP	TN	Loss
AE	0.808	0.958	0.800	52	545	1190	70443	0.0181
FC	0.817	0.961	0.813	65	501	1234	70430	0.0025
GAT	0.826	0.963	0.815	78	489	1246	70417	0.0147
GCN	0.818	0.959	0.798	95	521	1214	70400	0.0273
GIN	0.813	0.960	0.806	71	517	1218	70424	0.0148
Graph-T-A	0.725	0.932	0.695	386	605	1130	70109	0.0024
Graph-T-Z	0.246	0.833	0.338	1759	1025	710	68736	0.0024
Graph-T	0.832	0.963	0.821	40	499	1236	70455	0.0024

An absolute error model (Graph-T-A) and an unsmoothed output (Graph-T-Z) were also considered, changing Equation 1 and removing Equation 2 respectively. The absolute error metric had larger error rates since there will be FPs from both sides of the distribution. This confirms that the error signal direction is important to consider in this task.

Removing the smoothing step at the output caused a large decrease in performance at test time. Average precision reduced from 0.832 to 0.246, a decrease of 70%. This would suggest that smoothing the outputs to minimise noise from predictions is also a critical step to improving the robustness of the anomaly detector.

4 CONCLUSION

In this study, a GNN was trained to predict the nominal processing conditions from the laser features and positional information represented as a graph. The output prediction was compared to new observations giving a measure of the difference between the two. We found that performing a smoothing step at the output was essential, while the method is robust when compared with a Gaussian distribution. In future work we will test the method on more complex geometries.

REFERENCES

- T DebRoy, T Mukherjee, HL Wei, JW Elmer, and JO Milewski. Metallurgy, mechanistic models and machine learning in metal printing. *Nature Reviews Materials*, 6(1):48–68, 2021.
- Austin Derrow-Pinion, Jennifer She, David Wong, Oliver Lange, Todd Hester, Luis Perez, Marc Nunkesser, Seongjae Lee, Xueying Guo, Brett Wiltshire, et al. Eta prediction with graph neural networks in google maps. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*, pp. 3767–3776, 2021.
- Pablo Gainza, Freyr Sverrisson, Frederico Monti, Emanuele Rodola, D Boscaini, MM Bronstein, and BE Correia. Deciphering interaction fingerprints from protein molecular surfaces using geometric deep learning. *Nature Methods*, 17(2):184–192, 2020.
- Paul A Hooper. Melt pool temperature and cooling rates in laser powder bed fusion. *Additive Manufacturing*, 22:548–559, 2018.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016.
- Sebastian Larsen and Paul A Hooper. Deep semi-supervised learning of dynamics for anomaly detection in laser powder bed fusion. *Journal of Intelligent Manufacturing*, 33(2):457–471, 2022.
- Alvaro Sanchez-Gonzalez, Jonathan Godwin, Tobias Pfaff, Rex Ying, Jure Leskovec, and Peter Battaglia. Learning to simulate complex physics with graph networks. In *International conference on machine learning*, pp. 8459–8468. PMLR, 2020.
- 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, 2018.
- Yunsheng Shi, Zhengjie Huang, Shikun Feng, Hui Zhong, Wenjin Wang, and Yu Sun. Masked label prediction: Unified message passing model for semi-supervised classification. *arXiv preprint arXiv:2009.03509*, 2020.
- Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. Graph attention networks. *arXiv preprint arXiv:1710.10903*, 2017.
- Tian Xie and Jeffrey C Grossman. Crystal graph convolutional neural networks for an accurate and interpretable prediction of material properties. *Physical review letters*, 120(14):145301, 2018.
- Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural networks? *arXiv preprint arXiv:1810.00826*, 2018.