COMPARING PRINCIPAL COMPONENT ANALYSIS (PCA) AND β -VARIATIONAL AUTOENCODER (β -VAE) FOR ANOMALY DETECTION IN SELECTIVE LASER MELTING (SLM) PROCESS DATA

J. VOIGT¹ AND M. MOECKEL²

¹University of Applied Sciences Aschaffenburg, ZeWiS
Glanzstoffstr. 1, 63784 Obernburg a. Main, GERMANY
jorrit.voigt@th-ab.de and www.th-ab.de

²University of Applied Sciences Aschaffenburg
Wuerzburger Str. 45, 63739 Aschaffenburg, GERMANY
michael.moeckel@th-de.de and www.th-ab.de

Key words: Anomaly detection, variational autoencoder, machine learning, process monitoring, additive manufacturing, selective laser melting (SLM)

Abstract. The usability of machine learning approaches for the development of in-situ process monitoring, automated anomaly detection and quality assurance for the selective laser melting (SLM) process receives currently increasing attention. For a given set of real machine data we compare two established methods, principal component analysis (PCA) and β -variational autoencoder (β -VAE), for their applicability in exploratory data analysis and anomaly detection. We introduce a PCA-based unsupervised feature extraction algorithm, which allows for root cause analysis of process anomalies. The β -VAE enables a slightly more compact dimensionality reduction; we consider it an option for automated process monitoring systems.

1 INTRODUCTION

Machine learning techniques have already entered the stage to analyse SLM process and sensor data. Shevchik et al. [1] presented a study using acoustic emission of the SLM process as an input into a spectral convolutional neural network (SCNN), which classified the quality of the printed layer with an accuracy between 83-89% for three quality categories. Grasso et al. [2] proposed a method to detect defects related to overheating by using the principal component analysis (PCA) on process images for defining a statistical descriptor and applying a k-means clustering on this descriptor subsequently. Uhlmann et al. [3] achieved an accuracy of 63% in classifying the outcome of process data by using a Bayes classifier and used k-means clustering for identifying the patterns for the given labels in the data. To improve process monitoring for the SLM process other authors developed specific tools for in situ process monitoring [4–6]. Qi et al. [7] pointing out the need for machine learning tools in the field of additive manufacturing (AM) and labelled data, which requires a deep cooperation between computer and materials scientists as well as machine operators. Additionally, they are emphasizing among others the

importance of a suited data preprocessing and quality assurance for every layer.

In this work a comparison of the principal component analysis (PCA) and β -variational autoencoder (β -VAE) for the purpose of anomaly detection in SLM process data is presented. A feature extraction algorithm (FEA) based on the PCA for anomaly detection and root cause analysis is introduced. Furthermore, an optimization strategy to detect anomalies by using the β -VAE is described.

2 METHODOLOGY

Both PCA and the β -VAE can be used for dimensionality reduction and feature extraction [8, 9]. The choice of suitable features, however, requires assessment metrics. In this work, we compare the performance of both methods for the purpose of anomaly detection, i.e. we relate the number of detected anomalies to the total number of anomalies in an expert-labelled dataset.

In PCA, principal components (PC) are constructed as eigenvectors of the covariance matrix of the original data. A standard measure for their relevance is by the corresponding eigenvalue, which represents the fraction of the total data variance explained by the PC. This notion of relevance, however, does not necessarily agree with their importance for anomaly detection, as we will illustrate shortly. Hence, meaningful feature extraction requires selecting those dimensions that are relevant for anomaly detection. For further insights into PCA we refer to the ample literature, e.g. [10–12].

The β -VAE represents input data in a latent space representation of reduced dimensionality. An encoder and a decoder unit implement parametrized stochastic mappings between input space and latent space as neural networks. Training is by optimizing data reconstruction after passage through both the encoder and decoder network. The most important tuneable parameters of the model are the dimensionality of the latent space and the hyperparameter β , which regularizes the relative importance of the reconstruction error and disentangling in the latent space representation. Further details can be found, for instance, in Kingma and Welling [13, 14]. Note that training is w.r.t. the reconstruction error, not w.r.t. labels. Subsequent anomaly detection schemes can operate as unsupervised learning. Labelled datasets allow assessing their performance. For details about the used machine learning classifiers One-Class Support Vector Machine (OC-SVM), Mahalanobis Distance (MD) and clustering algorithm Density Based Spatial Clustering for Applications with Noise (DBSCAN) the reader is referred to [11, 15–17].

3 DATASET

The used data set originates from an experimental research group working on a DMG Mori Lasertec 30 SLM machine and consists of 38 building jobs with various geometries and variable duration between 200 and over 1600 layers. Time series data of 46 sensors has been recorded but split into independent time slices for each layer. As process parameters typically develop on comparatively long-time scales, aggregated sensor data at ten time points per printed layer has been defined and analysed in a layer wise approach. Labels indicating anomalies have been assigned to each layer based on the opinion of the machine operators and knowledge of materials scientists. The whole data set comprises 22215 layers with 95 anomalies from 38 building jobs. For algorithm training, we split the data into 60% training, 20% validation and 20% test set.

4 RESULTS AND DISCUSSION

4.1 PCA for dimensionality reduction and anomaly detection by classification

We investigated the performance of anomaly detection algorithms acting on latent spaces of variable dimensions. A truncated set of ordered principal components (PC) spans the latent space. By extending the set, the fraction of the total variance explained by the latent space increases. Starting from the first two PCs, which capture approximately 28% of the total variance, we observe a clustering of single layer data points w.r.t. building jobs but no clear separation of anomalies (compare Figure 1 (a)). This corresponds to our expectations that different geometries and building patterns dominate the variance of the data set. The marked ‘Cluster 1’ comprises of layers varying in geometry and typical process parameter spreading from the other layers. It represents exceptional layers with intended parameter modifications.

However, we observe that a clear detection of process anomalies is possible in latent spaces that include other than just the leading PCs. The F_1 scores for the three different anomaly detection algorithms applied on the latent spaces capturing at least 28% to 99.99% of the total variance are shown in Figure 1 (b). The scores are in between 0.15-0.25 on the first two PCs for all the classifiers. Adding dimensions to the latent space increases the F_1 score to a range between 0.6-0.65 for the DBSCAN and MD respectively and to 0.4-0.6 for the OC-SVM until 75% of the total variance. Increasing the captured variance to at least 80% (17-dimensional space) improves the performance for the DBSCAN and MD to a range between 0.7-0.77, while the OC-SVM increases the performance continuously from 0.62 at 81% captured variance to 0.75 at 99.99% of the total variance. In Figure 1 (c) a subspace determined by FEA for separating regular process layers and anomalies is shown. For further reference we single out ‘Cluster 2’, which represent a characteristic malfunction. For detailed discussion see section 4.2.

Figure 1: Results for the PCA are shown for 18 building jobs: (a) the visualization in the first 2 PCs, the samples are coloured and dimensioned according to their building process number, (b) the F_1 scores in relation to the captured variance and dimension, (c) an optimal subspace for anomaly detection determined by the FEA.

We understand that the first two PCs represent the variance coming from different building processes. Due to the flexibility of the AM process, each building process is unique and varies from others. Detection of process anomalies is substantially improved by later PCs, as can be seen by the increasing F_1 score with the dimension of the latent space (i.e. captured variance). The significant changes in the performance of the classifiers by adding PC 3 and the PCs 14-

17 to the investigated latent space demonstrates that the anomalies are represented in certain PCs and are distinguishable from statistical noise as well as from the variance due to the process.

However, adding PCs to the investigated latent space increases the dimension which potentially leads to higher runtimes for detection algorithms. Furthermore, the captured variance does not indicate the relevance of a PC for anomaly detection and the resulting high dimensional subspace makes the detection process incomprehensible, because the required time for investigating all eigenvector loadings is prohibitive. Addressing these two issues, a feature extraction algorithm based on the PCA for anomaly detection and root cause analysis is presented in section 4.2.

4.2 Feature extraction algorithm and root cause analysis

We present a quick PCA based feature extraction algorithm (FEA) for dimensionality reduction and root cause analysis. It rests on two global parameters, the detection threshold (*detect_threshold*) for outlier detection in each PC and the global detection percentile (*glob_pctl*) and two optional parameters for the desired captured variance (*cap_var*) and resulting dimension (*m*). The pseudocode is given in Algorithm 1. The FEA takes a z-score standardized dataset matrix *D* as input and outputs a set of PCs, which are spanning the subspace with an optimal separation between anomalous and normal data.

Algorithm 1: Feature Extraction Algorithm (*var_cap, detect_threshold, glob_pctl, m*)

Input: A z-standardized dataset matrix *D*

Output: A set *S* of *m* principal components

1. $(X, \lambda) \leftarrow \text{transformed representation and ordered eigenvalues by PCA}(D)$
2. $n \leftarrow \min_n (\sum_1^n \lambda_i \geq \text{var}_{cap})$
3. For *j* in 1:*n*
4. $X_j \leftarrow \text{PC}_j(X) = X[:, j]$
5. $\sigma_{X_j} \leftarrow \sqrt{\text{Var}(X_j)}$
6. $\tilde{X}_j \leftarrow X_j \div \sigma_{X_j}$
7. $\text{score}_{local}^j \leftarrow \tilde{X}_j \geq \text{detect_threshold}$
8. $d_M(\tilde{X}) \leftarrow \sum_{j=1}^n \tilde{X}[:, j]$
9. $\text{score}_{global} \leftarrow \{x \text{ in } \tilde{X} | \Phi(d_M(\tilde{X})) \geq \text{glob_pctl}\}$, where Φ is CDF
10. $J = [1, \dots, n]$; $A = [\text{score}_{global}]$; $S = []$
11. For *i* in 1:*m*
12. For *j* in *J*
13. $F_1^j \leftarrow F_1(A, \text{score}_{local}^j)$
14. $t \leftarrow \max_j (\{F_1^j\}_{j \in J})$
15. $J = J.\text{remove}(t)$
16. $S = S.\text{append}(t)$
17. $A = A.\text{remove}(\text{score}_{local}^t)$
18. Return *S*

At first, a global score for anomalies is determined by iterating through the 1-dimensional representations in the transformed space of the PCs. To obtain local and global scores for anomaly detection, the transformed data in each selected PC (X_j) is rescaled w.r.t. its standard deviation (σ_j) to obtain an anomalousness score ($\tilde{X}$) for each sample x_j^i . By adding the anomalousness scores from each latent space (step 8) the Manhattan distance (d_M) for each sample is calculated. A global score for anomalies is obtained by applying an introduced percentile ($glob_{pct}$) e.g., 0.995. Simultaneously the samples x_j^i in each representation are distinguished in normal data and outliers by a threshold to receive local candidates for anomalies in every PC (step 7). In step 10, three lists are initialized. Here J is the list of considered PCs, A is the list of anomalies defined by the global score and S the list of selected PCs. Then the F_1 scores (or any other valid metric) from the global scores (list A) and the local scores are computed. The PC $_i$ ($i \in J$) with the highest F_1 result is selected. For further selection of PCs, the prior determined PC is deleted from list J and those anomalies, represented by the local score, are removed from list A . Hence, subsequently selected PCs represent further anomalies.

The determined PCs using the introduced FEA on the investigated data set are PC 16 followed by PC 4 and PC 20. The reached F_1 metric on the 3-dimensional space is 0.77. The latent space spanned by the first two determined PCs (16 and 4) is presented in Figure 1 (c) and a separation of normal and anomalous data points is observable.

Beside of detecting anomalies, the FEA can be used for root cause analysis. The PCs are weighted linear combinations of the attributes. This allows for observation of the most influential attributes on each PC and to take their covariances into account. Investigating the loadings of the selected PCs provides the opportunity to trace back the error signals and the malfunction in the system. The loadings from PC 4 and PC 16 are shown in Figure 2, in which the red squares indicate the most significant loadings. In PC 16 the two most weighted signals belong to a sensor from the powder feeding system and the sonic sieve state for powder processing. In the 4th PC two significant loadings are the signals from oxygen measurements and the other two are monitoring the valve of the inert gas system.

Register for free at <https://www.scipedia.com> to download the version without the watermark

Figure 2: Weights of the eigenvectors from the determined PC 4 (a) and PC 16 (b).

The root cause analysis demonstrates that PC 16 and PC 4 represent two types of anomalies in the data set. The PC 16 shows issues with the powder processing system, whereas PC 4 reveals malfunction in the process atmosphere. The emergence of different failure classes is the consequence of the removal of local anomalies in step 17 (compare Algorithm 1), which forms disjunct subsets of anomalies. The marked ‘Cluster 2’ belongs to a malfunction in the process atmosphere and is primarily represented by PC 4.

4.3 β -variational autoencoder for automatic anomaly detection

A β -VAE consists of an encoder and decoder unit, which map data between the input and the latent space representation. Training is w.r.t. a cost function formed of a reconstruction error and a regularization term representing the Kullback-Leibler divergence. The relative weight of the regularization is determined by β . The latent space is fixed at dimension $d = 2$. The network is trained on unclean data containing anomalies. To separate off anomalies in the 2-dimensional latent space, the β -value is adapted and three different standard classifiers (DBSCAN, OC-SVM, MD) have been used. We evaluate the approach by calculating the F_1 score for detected anomalies w.r.t. all known anomalies in the labelled dataset.

The reached F_1 scores for different β -values are presented in Figure 3 (a). At $\beta = 1$ there is no significant detection of anomalies ($F_1 = 0.15$) for all three classifiers. By decreasing the β -value to 10^{-1} the F_1 metrics increase to a range between 0.5-0.55 and beyond the threshold $\beta \leq 10^{-2}$ the F_1 scores are between 0.55-0.71, only depending on the classifier used. Figure 3 (b) shows the separation of anomalies at threshold $\beta = 10^{-2}$ and a clustering marked as ‘Cluster 2’, which comprises the same data points as the ‘Cluster 2’ obtained from the subspace of PC 4 and PC 16 above. Zooming into the origin shows separation of the different building jobs as illustrated in Figure 3 (c). The ‘Cluster 1’ represents the same data points as ‘Cluster 1’ from the leading PCs (cf. section 4.1).

Figure 3: Results for the β -VAE: (a) the F_1 scores for the classifiers in relation to the β -value, (b) the visualized embedding at $\beta = 10^{-2}$ with a clear separation of normal and anomalous samples, (c) the disentangling of the different building jobs, where the samples are coloured and dimensioned according to their process number.

On the one hand the smaller β -parameter leads to a less weighted Kullback-Leibler divergence in the loss function and therefore to a less constraint embedding, which improves disentangled and information conserving encodings of the data. On the other hand the decoder samples from this embedding to reconstruct the input data. A clear distinction of most anomalies is possible in latent space for β below threshold. This is achieved as the reconstruction loss dominates the objective function. The remaining regularization confines the distribution in latent space such that informative graphical representations can be obtained.

We identify most anomalies which form ‘Cluster 2’ in Figure 1 (c) and observe that these anomalies form a well-located cluster in latent space of the β -VAE. Similarly, we confirmed the localization of elements of ‘Cluster 1’ in latent space. The ‘Cluster 1’ represents the variance

due to geometry and typical process parameter spreading, whereas ‘Cluster 2’ is formed by one class of malfunctions in the gas system. These results are similar to the findings of Alemi et al. [18], who showed the separation of the MNIST images and clustering of same or similar looking numbers on adapted β -parameters and to the findings of Higgins et al. [19] who presented an automatic unsupervised routine for interpretable image data representation.

The results in this section demonstrate that the encoding of the neural network can be used for anomaly detection and effective information conserving dimensionality reduction by adapting the β -parameter below a certain threshold. On different scales of the latent space variables two contributions to the total variances can be separated: close to origin the variance due to the building jobs is visible, outside of a β -dependent boundary around the origin process anomalies can be found. This makes the β -VAE a promising technique for automatic anomaly detection and exploratory data analysis (EDA) in process monitoring.

5 COMPARISON

We compare the two investigated methods with each other regarding their performance for anomaly detection and exploratory data analysis. If PCA is applied for representing data as a truncated expansion in terms of PCs with respect to the captured variance, the F_1 score saturates roughly for a 17-dimensional space. This naïve conclusion on the required size of the latent space, however, does not account for the particular nature of the discussed anomalies.

However, applying Algorithm 1, the introduced feature extraction algorithm (FEA), and demanding $F_1 \geq 0.7$, results in a 3-dimensional latent space. This massive reduction of latent space dimensionality is possible as FEA implements a further description of anomalies based on two predefined parameters: *detect threshold* for defining local scores within each PC as well as the *glob_{pc1}* for global scoring. It facilitates performing a root cause analysis on the small number of selected PCs.

The β -VAE showed a comparable F_1 for only 2 dimensions. Anomaly detection could be achieved for $\beta \leq \beta_{th}$, i.e. below a certain threshold value $\beta_{th} \approx 0.01$. Small but nonzero β -values allow for informative visualization of total variance, which may support exploratory data analysis. Table 1 summarizes the results from these different methods.

Table 1: Summary of the used methods for reducing the dimension and anomaly detection. The given F_1 scores are those reached at the shown dimension.

Technique	Metrics			
	F_1	Root cause analysis	Exploratory data analysis	Dimension
PCA	0.75	0	0	17
PCA & FEA	0.77	+	0	3
β -VAE	0.71	—	+	2

‘+’ = Advantage | ‘0’ = fair | ‘—’ = Not possible

All three approaches agree on clustering exceptional layers (‘Cluster 1’) and particular anomalies (‘Cluster 2’). However, root cause analysis is only possible by using PCA based methods.

6 CONCLUSION AND OUTLOOK

We investigated the applicability of PCA and β -VAE for the purpose of dimensionality reduction and unsupervised anomaly detection in machine process data. The time-series data was obtained from 38 SLM printing processes, for which process anomalies have been labelled by experts. We perform exploratory data analysis and compare three approaches for semi-automatic feature extraction. Each of them comprises of a method for dimensionality reduction and a subsequent standard classifier for anomaly detection. From this comparison the following conclusions can be drawn:

- Anomalies can be detected by all three techniques, which, however, require for comparative performance latent spaces of different dimensions.
- A straightforward approach by truncating an expansion in principal components following the objective to capture a pre-defined amount of explained variance with a minimum of PC, leads to unnecessary high dimensional subspaces.
- Massive reduction can be achieved by selecting latent spaces with respect to anomaly characteristics. For this purpose, a specific Feature Extraction Algorithm (FEA) based on PCA has been introduced.
- FEA provides the opportunity for root cause analysis by inspecting selected principal components.
- We use a β -VAE to project the data into a 2-dimensional latent space. The degree of separation between anomalies and regular data can be tuned by the β -parameter, which allows for exploratory data analysis, visualization, and anomaly detection. For the sole purpose of anomaly detection, no significant advantages have been found at $\beta > 0$.
- Data points located within a boundary around the origin of latent space represent data of regular building processes. Anomalies can be found outside. For values of β below a certain threshold ($\beta < \beta_{th}$) the separation of anomalies measured by F1 score gets independent of β .
- Both FEA and β -VAE detect and cluster anomalies in a similar way. Clusters formed by FEA can be re-discovered as clusters in latent space of the β -VAE.

We consider FEA and β -VAE both candidates for semi-automatic and unsupervised anomaly detection. We suggest further studies to evaluate the computational efficiency of the presented methods and to include further anomaly detection techniques e.g., by measuring the reconstruction error of an autoencoder trained on practically anomaly-free data.

REFERENCES

1. Shevchik SA, Kenel C, Leinenbach C et al. (2018) Acoustic emission for in situ quality monitoring in additive manufacturing using spectral convolutional neural networks. Additive Manufacturing 21:598–604.
2. Grasso M, Laguzza V, Semeraro Q et al. (2017) In-Process Monitoring of Selective Laser Melting: Spatial Detection of Defects Via Image Data Analysis. Journal of Manufacturing Science and Engineering 139:255.

3. Uhlmann E, Pontes RP, Laghmouchi A et al. (2017) Intelligent Pattern Recognition of a SLM Machine Process and Sensor Data. *Procedia CIRP* 62:464–469.
4. Craeghs T, Bechmann F, Berumen S et al. (2010) Feedback control of Layerwise Laser Melting using optical sensors. *Physics Procedia* 5:505–514.
5. Kolb T, Gebhardt P, Schmidt O et al. (2018) Melt pool monitoring for laser beam melting of metals: assistance for material qualification for the stainless steel 1.4057. *Procedia CIRP* 74:116–121.
6. Krauss H, Zeugner T, Zaeh MF (2014) Layerwise Monitoring of the Selective Laser Melting Process by Thermography. *Physics Procedia* 56:64–71.
7. Qi X, Chen G, Li Y et al. (2019) Applying Neural-Network-Based Machine Learning to Additive Manufacturing: Current Applications, Challenges, and Future Perspectives. *Engineering* 5:721–729.
8. Kuncheva LI, Faithfull WJ (2014) PCA feature extraction for change detection in multi-dimensional unlabeled data. *IEEE Trans Neural Netw Learn Syst* 25:69–80.
9. Nugroho H, Susanty M, Irawan A et al. (2020) Fully Convolutional Variational Autoencoder For Feature Extraction Of Fire Detection System. *Journal of Computer Sciences and Information* 13:9.
10. Jolliffe IT (2002) *Principal Component Analysis*, Second Edition. Springer Series in Statistics. Springer-Verlag New York Inc, New York, NY.
11. Zaki MJ, Meira W (2014) *Data mining and analysis: Fundamental concepts and algorithms*, 1st. Cambridge University Press.
12. Bishop CM (2016) *Pattern Recognition and Machine Learning*, Softcover reprint of the original 1st edition 2006 (corrected at 8th printing 2009). Information science and statistics. Springer New York, New York, NY.
13. Kingma DP, Welling M (2013) Auto-Encoding Variational Bayes.
14. Kingma DP, Welling M (2019) An Introduction to Variational Autoencoders. *FNT in Machine Learning* 12:307–392.
15. Sander J, Elp M, Krusch H, Paterl (1998) Density-Based Clustering with Spatial Data bases: The Algorithm GDBSCAN and Its Applications. *Data Mining and Knowledge Discovery* 2:169–194.
16. Schölkopf B, Williamson R, Smola A et al. Support Vector Method for Novelty Detection. In: *Advances in Neural Information Processing Systems* 12, vol 12, pp 582–588.
17. Simoudis E, Han J, Fayyad UM (eds) (1996) *Proceedings / Second International Conference on Knowledge Discovery & Data Mining*. AAAI Press, Menlo Park, Calif.
18. Alemi AA, Fischer I, Dillon JV et al. (2017) Deep Variational Information Bottleneck. *ICLR*.
19. Irina Higgins, Loïc Matthey, Arka Pal et al. (2017) beta-VAE: Learning Basic Visual Concepts with a Constrained Variational Framework. *ICLR*.

Register for free at <https://www.scipedia.com> to download the version without the watermark