


A novel machine learning system for early defect detection in 3D printing

Michał Błaszczkowski^{1*} , Dariusz Majerek¹, Elżbieta Sędziewska¹,
Paweł Tomiło², Jarosław Pytka³

¹ Faculty of Mathematics and Information Technology, Lublin University of Technology, Nadbystrzycka 38, 20-618 Lublin, Poland

² Faculty of Management, Lublin University of Technology, Nadbystrzycka 38, 20-618 Lublin, Poland

³ Faculty of Mechanical Engineering, Lublin University of Technology, Nadbystrzycka 36, 20-618 Lublin, Poland

* Corresponding author's e-mail: m.blaszczkowski@pollub.pl

ABSTRACT

This paper discusses a comprehensive study to develop a machine learning model for detecting unwanted vibrations during the 3D printing process. Undesired vibrations can significantly degrade print quality, leading to defects such as void formation, poor surface quality and improper layer bonding. Identifying and mitigating these vibrations is essential to ensuring the reliability and precision of 3D printed products, which is particularly crucial in sectors such as healthcare, automotive, and aerospace. The study introduced a novel system with an inertial measurement unit (IMU) mounted on the printer head, which records acceleration and angular velocity in three axes. The data is transmitted to a microcontroller and then to an acquisition device that controls a controlled vibration generator. The collected information formed a dataset for training and testing various machine learning models. Of all the models evaluated, the Dense Neural Network (DNN) showed the highest performance in accurately distinguishing normal print vibrations from unwanted vibrations. The study underscores the critical importance of early defect detection, which saves time and reduces costs, being essential for the widespread adoption of incremental manufacturing technology. Early identification of defects enables immediate intervention and correction of errors before they become serious defects affecting the quality of the final product. This is particularly important in the context of increasing automation and optimization of manufacturing processes.

Keywords: 3D printing, extraneous vibrations, machine learning.

INTRODUCTION

3D printing technology is a popular additive manufacturing method/technology that allows the creation of a previously designed object through methods such as stereolithography (SLA) or deposition modeling (FDM). This fast-emerging technology creates objects from their geometrical representation by successive addition of materials and is increasingly used for mass customization, production of any type of open-source design in the field of agriculture, in healthcare, automotive industry, locomotive industry and aviation industries [1]. The popularity of this method makes that several different studies have been conducted,

including different materials, multi-materials, advanced methodologies, design, and possible optimization [1–4]. According to Parvanda et al. [5], additive manufacturing technologies have numerous benefits and advantages, they also have some limitations and challenges that hinder their complete adoption, among which we can point unavoidable failures and errors due to limited experience and knowledge, also related to different materials used or combined with. Another important research challenge follows from the need to properly understand defects in 3D printing processes. The most common defects observed in FDM technology while, for example, depositing fiber-reinforced composites are void formation, poor surface finish,

and improper bonding between the layers. Some defects are also associated with SLM, including the generation of residual stress and cracking [5].

Among the limitations listed above, process stability, defect identification, and elimination are considered key research challenges in 3D printing technologies. According to [5] the early identification of the defect can save a lot of time and cost. Integration of both hardware and physical systems is required for the same. To solve these problems, several researchers apply advanced image processing, or laser profilometers, to identify defects. Despite these new solutions and improvements, the adoption of additive manufacturing in the future, especially the usage of machine learning models to automatically detect problems during the 3D printing process. Algorithms of this kind can help detect unwanted vibrations that adversely affect the final print, as well as using image analysis techniques to detect defects and correct them at printing time [6].

The methods outlined above, which are used in presented 3D printing research, because of the low cost of the equipment, wide availability, and frequent use in nonprofessional applications. In each of the cases described, there is a movement in the axis of the printer head or worktable. Due to the wide availability of 3D printers, their location does not always allow one to eliminate the impact of vibrations, which are not related to the technology itself, but are a consequence of the process occurring around the printer. This mainly concerns personal applications of 3D printing.

The article identifies the occurrence and impact of vibrations not originating from the 3D printing process on the quality of the final print. In the remainder of the article, the term extraneous vibrations are adopted for the description of this kind of vibrations.

Vibrations that are generated by the printing process itself are related to the dynamic characteristics of the printer system, affect the quality of the final product, and more specifically, the surface roughness and mechanical properties of the print [7, 8]. In addition, extraneous vibrations can occur during the printing process, which can negatively affect the final product as well as the printing process itself. The effect of extraneous vibrations on the quality of the final 3D print, is not a topic that has been addressed in the literature.

The main objective of the research described in this article is to create a model based on machine learning methods to detect the occurrence of extraneous vibrations. For this purpose, an IMU system was mounted on the head of the 3D printer, which

measures the acceleration and angular velocity in three axes. In order to control the occurrence of extraneous vibrations, a vibration generator was developed. Such a setup made it possible to check the transmission of extraneous vibrations to the printer head, and by being able to control the occurrence of extraneous vibrations, it was ensured that the moment of their occurrence and their duration were known. This allowed the creation of a data set for machine learning models, the task of which will be to detect extraneous vibrations.

MATERIALS AND METHODS

Measuring station

To study the previously described above, it was decided to use a 3D printer that works using FDM technology with a moving head on the z-axis. An adequate test stand was designed and constructed [9]. An IMU was mounted on the printer head to transmit magnetometer and gyroscope readings to microcontrollers and further to a data acquisition device. The data acquisition device is also the controller of the vibration generator. The vibration generator was designed as a mechanical device based on the Scotch Yoke mechanism [10]. A schematic diagram of the test stand is shown in Figure 1.

The measurements were carried out in such a way that a vibration generator was operated at 15 s intervals for 5 s throughout the printing period. The time periods in which vibrations occurred were recorded, making it possible to create a data set for classification. A photo of the test stand is shown in Figure 2.

Data preprocessing and machine learning models

In order to select a detect extraneous vibration model, the following models were checked [11]: LGBM, Gradient Boosting Classifier (GB), Random Forest Classifier (RF), Extra Trees Classifier (ET), Decision Tree Classifier (DT), Quadratic, Discriminant Analysis (QDA), Ada Boost Classifier (AB), Naive Bayes (NB), Support Vector Machine (SVM), K Nearest Neighbours Classifier (KNN), Logistic Regression (LR), Ridge Regression Classifier (RR) and Dense Neural Network (DNN). The dataset consists of 2 classes, where class 0 indicates vibrations related to the printing

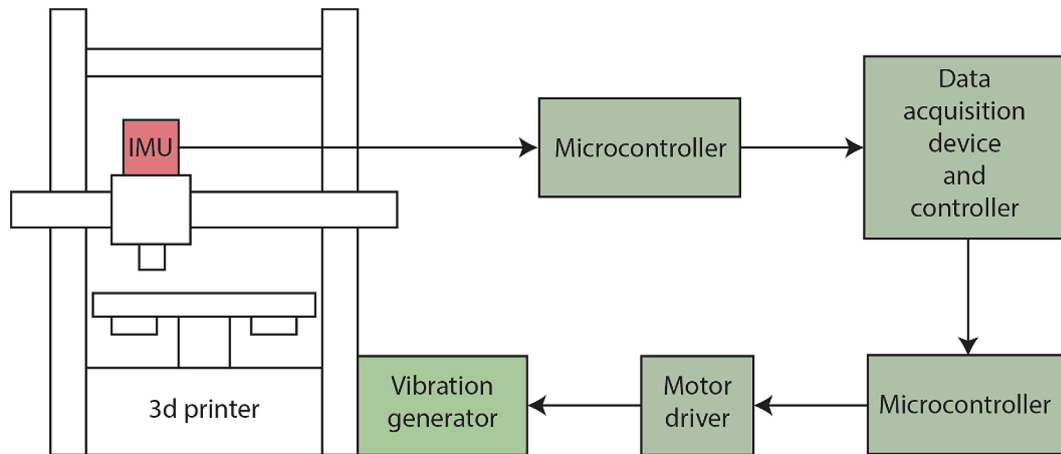


Figure 1. Schematic of the test stand




Figure 2. Photo of the test stand

process and class 1 indicates the occurrence of extraneous vibrations. The verification of printed specimens was conducted by a 3D printing expert, focusing on surface quality, absence of void formations, and proper layer bonding. A total of six test prints were produced under varying conditions: two without any extraneous vibrations, two with environmental vibrations, and two with vibrations induced by a generator. These conditions allowed for comparative analysis of the impact of vibrations on print quality.

In addition, a baseline model (Dummy Classifier, DC) was created that classified all samples as correct values in class 0. Using a baseline model in machine learning provides a simple comparison point for more complex models, ensuring that the improvements offered by advanced algorithms are meaningful and worth the additional complexity and computational cost. Baseline models, often simple and straightforward, help to understand the basic performance threshold any sophisticated model should exceed to be considered valuable. This approach prevents overengineering and helps in identifying whether the complexity of a new model genuinely translates to better performance, guiding researchers and practitioners towards more effective and efficient machine learning solutions.

Due to the system's design and the assumptions made, the resulting data contained three times more observations assigned to the class indicating the absence of unwanted vibrations compared to the class indicating their occurrence during the 3D printing process. The problem of imbalanced classes poses a significant challenge in data analysis, as machine learning models tend to perform better at recognizing the majority class. Lower effectiveness in identifying the minority class negatively impacts the overall quality of the analysis, which can lead to erroneous conclusions and limit the practical applicability of the results in real-world scenarios. For this reason, certain techniques were employed to effectively balance the dataset while preserving the overall structure of each class.

The dataset was divided into a training set (70%) and a testing set (30%), with the proportions of each class preserved. Ultimately, the training

set contained 30175 observations for class 0 and 10058 observations for class 1, with each observation represented by six predictors. In the first step, for the majority class, the potential for creating a representative sample with a reduced number of observations was analyzed, utilizing samples of 20000, 15000, and 10000 observations. As a result, based on the analysis of density plots and the outcomes of Kolmogorov-Smirnov tests [12] comparing the distributions of individual predictors, it was determined that the sample containing 15000 observations is fully representative and can be effectively utilized in further analyses. Figure 3 shows that the density plots for the fifth predictor, both for the full population and the selected sample, nearly overlap, with only minimal deviations observed. In this case, the Kolmogorov-Smirnov test statistic was 0.0027, while the p-value was approximately equal to 0.9999, which clearly indicates that there are no significant differences between the distributions being analyzed. For the remaining predictors, the results were similarly consistent. Consequently, the sample of 15000 observations was selected for use in further analyses. The operation performed on the dataset enables a reduction in computation time and the saving of computational resources.

After selecting the representative sample for the class representing the absence of vibrations, the Synthetic Minority Over-Sampling Technique (SMOTE) [13] technique was employed to fully balance the distribution of both categories in the training set. SMOTE algorithm is used to address the issue of imbalance in datasets, where one

class significantly outnumbers another. It generates synthetic examples of the minority class by interpolating between existing minority class instances. This approach not only increases the quantity of minority class samples but also helps in diversifying the dataset. By enhancing the representation of the minority class, SMOTE aids in improving the performance of machine learning models on imbalanced data, leading to more accurate and fair predictions.

Results and their description

Based on the modifications previously conducted, the final training dataset was balanced to include 15000 observations for each category, which served as the foundation for constructing the machine learning models. The test set retained the original class proportions, with 12932 observations representing the absence of introduced vibrations (class 0) and 4311 observations representing external vibrations (class 1). To ensure optimal evaluation of the models, the following metrics were utilized: accuracy (A), AUC, recall (R), precision (P), Kappa, F1-score, and MCC. The explanation and description of these metrics can be found in the referenced publication [14]. The metric values on the test set for the models with adjusted probability thresholds are presented in Table 1.

Considering all the evaluated metrics, the DNN emerged as the best model. However, the other models also demonstrated a strong ability to model the data, achieving more satisfactory

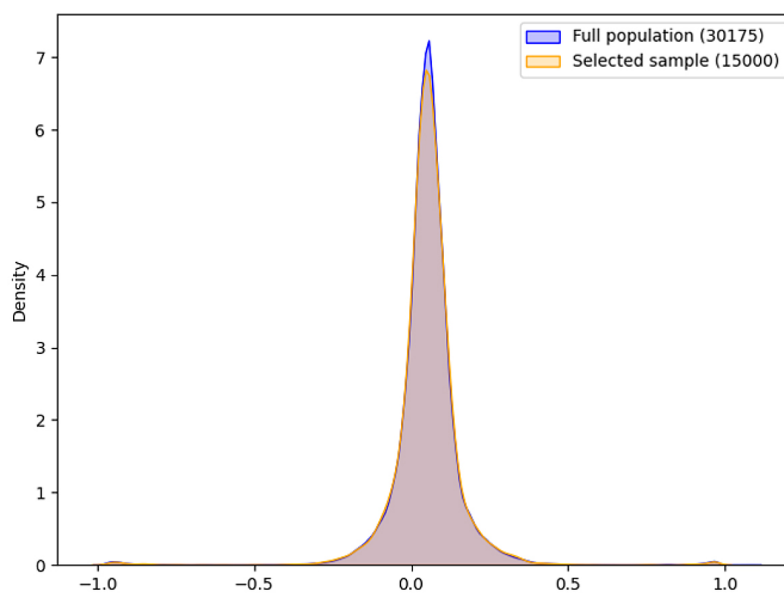


Figure 3. Comparison of density plots

Table 1. Metrics for the models studied

| Model | A | AUC | R | P | F1 | Kappa | MCC |
|-------|------|------|------|------|------|-------|------|
| DNN | 0.82 | 0.80 | 0.75 | 0.60 | 0.67 | 0.55 | 0.55 |
| LGBM | 0.70 | 0.73 | 0.63 | 0.44 | 0.51 | 0.30 | 0.32 |
| GB | 0.68 | 0.73 | 0.68 | 0.41 | 0.51 | 0.30 | 0.32 |
| RF | 0.65 | 0.75 | 0.74 | 0.40 | 0.52 | 0.28 | 0.32 |
| ET | 0.71 | 0.74 | 0.60 | 0.44 | 0.51 | 0.32 | 0.32 |
| DT | 0.64 | 0.60 | 0.55 | 0.35 | 0.43 | 0.19 | 0.19 |
| QDA | 0.72 | 0.66 | 0.36 | 0.42 | 0.39 | 0.20 | 0.20 |
| AB | 0.71 | 0.74 | 0.60 | 0.45 | 0.51 | 0.32 | 0.32 |
| NB | 0.71 | 0.65 | 0.35 | 0.40 | 0.38 | 0.19 | 0.19 |
| KNN | 0.66 | 0.70 | 0.63 | 0.40 | 0.49 | 0.26 | 0.28 |
| SVM | 0.71 | 0.73 | 0.61 | 0.44 | 0.51 | 0.32 | 0.32 |
| LR | 0.52 | 0.50 | 0.50 | 0.26 | 0.34 | 0.02 | 0.02 |
| RR | 0.52 | 0.00 | 0.50 | 0.25 | 0.34 | 0.02 | 0.02 |
| DC | 0.75 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 |

results than the baseline Dummy Classifier. A Dense Neural Network is a type of artificial neural network where each neuron is fully connected to every neuron in the subsequent layer. DNNs are widely used for various machine learning tasks such as classification, regression, and pattern recognition due to their flexibility and ability to model complex, non-linear relationships in data. The architecture of a DNN typically consists of an input layer, multiple hidden layers, and an output layer, with each layer composed of neurons that apply a weighted sum followed by an activation function. Training a DNN involves optimizing the weights through backpropagation, often using techniques such as stochastic gradient descent and various regularization methods to prevent overfitting and improve generalization. DNNs are particularly powerful in capturing intricate patterns in data and are becoming increasingly popular due to their effectiveness in a wide range of applications and their ability to handle complex datasets [15].

Hyperparameter tuning is a critical step in the machine learning model development process, aimed at optimizing the parameters that govern the model’s learning process. Unlike model parameters, which are learned automatically during training, hyperparameters are set prior to training and can significantly impact model performance. Techniques such as grid search, random search, and Bayesian optimization are commonly used for hyperparameter tuning, enabling the selection of optimal hyperparameter values that enhance the model’s ability to generalize to new data, thus

improving accuracy, reducing overfitting, and enhancing overall model effectiveness. For all algorithms, a hyperparameter tuning process was conducted to find the optimal parameters that maximize predictive performance. The tuned parameters were carefully selected based on their importance and relevance to each specific algorithm, ensuring that the models achieved the best possible results. In this paper, hyperparameters related to the learning process for the neural network, such as learning rate, number of epochs, and batch size, were tuned. The final values were set to a learning rate of 0.0001, 20 epochs, and a batch size of 32. The optimal architecture was determined based on a series of tests, resulting in a network consisting of three hidden layers with 256, 128, and 32 neurons, respectively. The RELU activation function, along with the ADAM optimizer, was applied to the neural network [16, 17].

Threshold optimization in machine learning involves fine-tuning the decision boundary that determines the classification of instances into different classes. This process is crucial in scenarios where classes are imbalanced or when the costs of different types of errors vary significantly. Adjusting the threshold, a model can be made more sensitive to either the positive or negative class. Effective threshold optimization can lead to improved model accuracy and better alignment with specific business objectives or operational requirements. The default decision threshold (0.5) was optimized by maximizing the F1-score metric and was adjusted to 0.53. It is worth noting that for most models, the predicted probabilities

were close to 0.5, indicating that they were not confident in their predictive decisions. This uncertainty could stem from the fact that the predictors were not sufficiently informative, and there were too few of them to allow the models to clearly distinguish between the classes. Only the dense neural network exhibited the majority of probabilities at extreme values. To enhance the system’s usability, an additional probability threshold was introduced based on the analysis. Specifically, if the probability of belonging to one of the two classes does not exceed 70%, the observation is not classified, and the measurement must be repeated. A lack of sufficient probability levels may suggest that certain issues occurred during measurements, either with the printer’s operation or its operational environment. Therefore, the developed system allows for the prior calibration of the device and adjustment of working conditions, thereby avoiding unnecessary and incorrect prints. It is more desirable to perform additional measurements and calibrations rather than producing defective 3D prints, considering the associated costs. Particularly as the system, over time and with the integration of new data, will undoubtedly become increasingly efficient. In our situation, 55% of all observations in the test set have been classified, and the detailed results for the final system are presented in the confusion matrix shown in Figure 4.

In summary, the model correctly classified 5880 instances from class 0, representing approximately 84% of all observations in that class, and 1705 instances from class 1, accounting for 75% of all observations in that group. A total of 1130 observations belonging to class 0 were incorrectly classified as class 1. However, this is not overly problematic, as in the developed early defect detection system, such observations will be re-evaluated for the presence of unwanted vibrations, which will not result in significant operational costs. In contrast, observations belonging to class 1 that were misclassified by the model as class 0 are significantly more troublesome. In such cases, no warning about external vibrations will be generated, which is likely to result in an incorrect print, leading to costs associated with wasted materials, time, and the need to redo the work. Nevertheless, such cases account for only 25%, with the remaining instances being correctly detected. Given the presented results, the percentage of incorrect predictions is relatively small. The system has the potential to significantly optimize the 3D printing process by effectively detecting unwanted vibrations that may affect the quality of the produced components. Additionally, an analysis of feature importance using shapley additive explanations (SHAP) will be conducted to gain a better understanding of the predictions made. SHAP is a method used to interpret machine learning models by attributing




Figure 4. Results for final system bases on DNN

the contribution of each feature to the final prediction [18]. It is based on cooperative game theory, specifically the shapley value, which ensures a fair distribution of contributions among features. SHAP values provide insight into how much each feature increases or decreases the predicted outcome. This method helps in understanding the inner workings of complex models by explaining individual predictions and identifying key features that drive decisions. SHAP is widely used for its consistency and accuracy in model interpretation across various types of machine learning models. The results of the SHAP analysis are presented in Figure 5.

The SHAP analysis indicates that the features *gz* and *ax* had the most significant impact on the model’s predictions (based on their ranking by the SHAP algorithm). Positive SHAP values indicate a higher probability of disturbances, while negative values suggest a lower chance of disturbances. The color gradient represents the value of each feature, with blue indicating low values and red indicating high values. The conclusions are not entirely clear-cut, as the groups of points often overlap. However, we can

observe a tendency where higher values of acceleration and angular velocity relative to specific axes generally suggest the presence of external vibrations and potential issues during the 3D printing process. A simultaneous increase in these parameters is typically indicative of disturbances in the system, suggesting that the machine may be experiencing external factors that could lead to issues in the printing process.

To effectively evaluate the effects of external vibrations on the 3D printing process, a series of tests were conducted in which three cubes were printed under different conditions. When the first cube was printed, there were no external vibrations of any kind. For the printing of the second cube, only environmental vibrations were present, while for the 3rd cube a vibration generator was used. Figure 6 shows the printouts of the cubes in question.

As can be seen from the figure above, there are no major differences in geometry between cubes numbered 1 and 2. In both cases, there is unevenness in the lower layers due to uneven cooling of the layers. In the case of cube 3, there are significant differences in geometry with




Figure 5. Feature importance based on SHAP values




Figure 6. Cubes prints. Red vertical sections indicate the width of the deformation in print

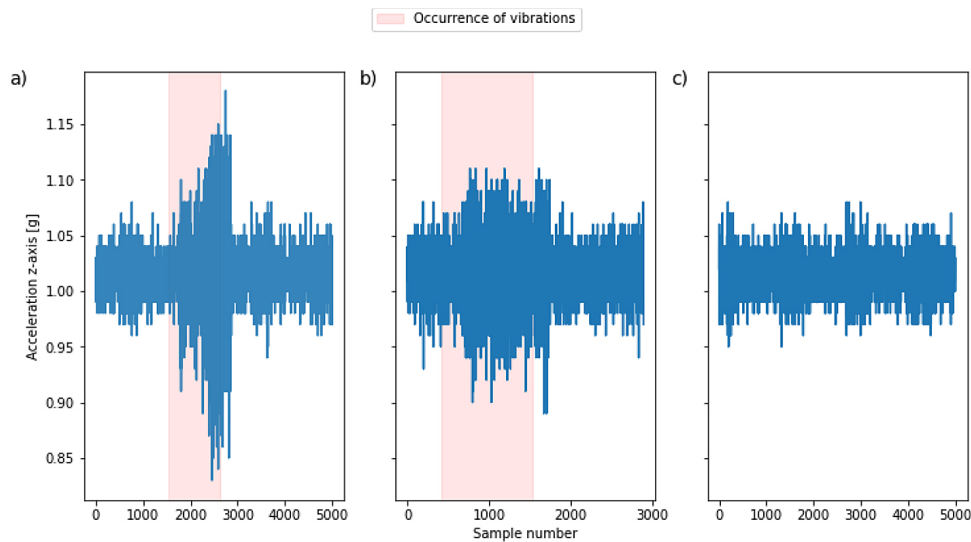


Figure 7. Example time series for: a) third cube, b) second cube, c) first cube

respect to cubes 2 and 1. A higher inequality in the initial layers is noticeable. Additionally, the entire surface is characterized by unevenly distributed bulges. For each of the cubes shown previously, an example of the acceleration run for the z-axis is shown. The results are presented only for the z-axis for the sake of clarity of the graphs and because on this axis the changes between the occurrence of vibrations are quite visible. For the third cube, it is noticeable that the amplitude of the vibration is increasing, even though the vibration generator was operating at the same frequency. For the 2nd cube, when checking the printing process manually, it is also possible to observe a change in the amplitude of vibrations. Figure 7 shows an example of time waveforms including the occurrence of vibration for 3 cubes.

CONCLUSIONS

The research described in the article successfully demonstrates the application of machine learning models, specifically the DNN, to detect extraneous vibrations in 3D printing processes. Through the integration of an IMU on the printer head, the study was able to generate a dataset that allowed the distinction between normal operational vibrations and detrimental extraneous vibrations. This approach not only highlighted the potential of machine learning in improving additive manufacturing processes but also underscored the challenges of dealing with

unbalanced data, which were effectively addressed through the selection of representative samples and the use of the Synthetic Minority Over-Sampling Technique (SMOTE) to enhance model performance.

Further analysis revealed that the DNN model outperformed the other evaluated models in accuracy and other key metrics, demonstrating its effectiveness in identifying extraneous vibrations with a refined balance of sensitivity and specificity. This was achieved through meticulous tuning of parameters related to the learning process and threshold optimization, underscoring the critical role of these steps in enhancing model accuracy and generalization to new data. The research highlights the importance of advanced analytical techniques in overcoming the limitations of traditional 3D printing processes, paving the way for more reliable and efficient manufacturing outcomes.

Extraneous vibrations affect the quality of the final print, their impact is mainly visible in the x and y axes since they cause the Z-wobbling phenomenon. They are transferred to the screws that move the head on the Y-axis, which translates into an increase in the occurrence of bulges form around the sides of a printed part.

According to the experiments, the occurrence of an extraneous vibrations of a small amplitude has little effect on the quality of the final print, but in the case of a larger amplitude, a resonance phenomenon may occur.

Due to the undefined origin of extraneous vibrations, the solution described in the article

can be utilized as a system responsible for detecting their occurrence and subsequently stopping the entire printing process until these vibrations cease. The machine learning model developed demonstrates the feasibility of detecting extraneous vibrations in 3D printing processes. However, this task presents significant challenges, primarily due to the unknown origins of such vibrations and their unpredictable occurrence. These vibrations are often caused by various processes occurring in the printer's environment, such as the operation of other machines. A key factor influencing the detection of these vibrations is the varying distance between the print head and the vibration source, which changes as the z-axis height adjusts during the printing process. As the print head moves further away from the source of disturbance, the transmission of vibrations diminishes, complicating the detection of extraneous vibrations when analyzing a single sample. One of the potential ways of detecting this type of vibrations is the use of an artificial neural network model, which will allow the analysis of data in windows to classify the signal characteristics (convolutional neural network) or by using Long-Short Term Memory (LSTM). To further enhance the system, it would be beneficial to expand the range of predictors related to the operational environment, allowing for even more accurate detection and classification of extraneous vibrations. To summarize, the presented system has significant potential in optimizing 3D printing processes by reducing material, time, and energy costs. By eliminating failed prints and ensuring optimal conditions before starting, it is possible to significantly reduce waste and accelerate the entire process, making the technology more efficient and reliable. Therefore, it is worth further developing and refining this system to fully realize its potential.

Acknowledgements

This work was prepared within the project “Innovative measurement technologies supported by digital data processing algorithms for improved processes and products”, contract number PM/SP/0063/2021/1, financed by the Ministry of Education and Science (Poland) as part of the Polish Metrology Programme.

REFERENCES

1. Shahrubudin, N., Lee, T.C., Ramlan, R. An overview on 3D printing technology: Technological, materials, and applications. *Procedia Manufacturing* 2019, 35, 1286–1296, <https://doi.org/10.1016/j.promfg.2019.06.089>
2. Nazir, A., Gokcekaya, O., Masum Billah, K., Ertugrul, O., Jiang, J., Sun, J., Hussain, S. Multi-material additive manufacturing: A systematic review of design, properties, applications, challenges, and 3D printing of materials and cellular metamaterials. *Materials & Design* 2023, 226, 111661, <https://doi.org/10.1016/j.matdes.2023.111661>
3. Prashar, G., Vasudev, H., Bhuddhi, D. Additive manufacturing: expanding 3D printing horizon in industry 4.0. *International Journal on Interactive Design and Manufacturing (IJIDeM)* 2023, 17, 2221–2235, <https://doi.org/10.1007/s12008-022-00956-4>
4. Hajare, D.M., Gajbhiye, T.S. Additive manufacturing (3D printing): Recent progress on advancement of materials and challenges. *Materials Today: Proceedings* 2022, 58, Part 2, 736–743, <https://doi.org/10.1016/j.matpr.2022.02.391>
5. Parvanda, R., Kala, P. Trends, opportunities, and challenges in the integration of the additive manufacturing with Industry 4.0. *Progress in Additive Manufacturing* 2023, 8, 587–614, <https://doi.org/10.1007/s40964-022-00351-1>
6. Brion, D.A.J., Pattinson, S.W. Generalisable 3D printing error detection and correction via multi-head neural networks. *Nat Commun* 2022, 13, 4654, <https://doi.org/10.1038/s41467-022-31985-y>
7. Kam, M., Saruhan, H., İpekçi, A. Investigation the effects of 3D printer system vibrations on mechanical properties of the printed products. *Sigma* 2018, 36(3), 655–66.
8. Kam, M., Saruhan, H., Ipekci, A. Investigation the effect of 3D printer system vibrations on surface roughness of the printed products. *DUBİTED* 2019, 7(2), 147–157, <https://doi.org/10.29130/dubited.441221>
9. Tomiło, P., Pytka, J., Jóźwik, J., Gnapowski, E., Muszyński, T., Łukaszewicz, A. Latent neural network for recognition of anomalies in 3D-print of a scale model for wind tunnel measurements. 2023 IEEE 10th International Workshop on Metrology for AeroSpace (MetroAeroSpace), Milan, Italy, 688–692, <https://doi.org/10.1109/MetroAeroSpace57412.2023.10190049>
10. Al-Hamood, A., Jamali, H., Oday, A., Senatore, A. Dynamics and lubrication analyses of scotch yoke mechanism. *International Journal on Interactive Design and Manufacturing (IJIDeM)* 2019, <https://doi.org/10.1007/s12008-019-00545-y>.
11. Hastie, T., Tibshirani, R., Friedman, J. Elements of

- Statistical Learning, Springer, 2009.
12. Dimitrova, D. S., Kaishev, V. K., Tan, S. Computing the Kolmogorov-Smirnov distribution when the underlying CDF is purely discrete, mixed, or continuous. *Journal of Statistical Software* 2020, 95(10), 1–42.
 13. Fernández, A., García, S., Galar, M., Prati, R. C., Krawczyk, B., Herrera, F. *Learning from Imbalanced Data Sets*. Springer 2018.
 14. Sokolova, M., Lapalme, G. A systematic analysis of performance measures for classification tasks. *Information Processing & Management* 2009, 45, 427–437, <https://doi.org/10.1016/j.ipm.2009.03.002>
 15. LeCun, Y., Bengio, Y., Hinton, G. Deep learning. *Nature* 2015, 521, 436–444, <https://doi.org/10.1038/nature14539>
 16. Glorot, X., Bordes, A., Bengio, Y. Deep sparse rectifier neural networks. *Proceedings of the 14th International Conference on Artificial Intelligence and Statistics (AISTATS)* 2011, 15, 315–323.
 17. Kingma, D. P., Ba, J. Adam: A method for stochastic optimization. *Proceedings of the 3rd International Conference on Learning Representations (ICLR)* 2015, <https://arxiv.org/abs/1412.6980>.
 18. Lundberg, S. M., Lee, S.-I. A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems* 2017, 30, 4765–4774. https://papers.nips.cc/paper_files/paper/2017/file/8a20a8621978632d76c43dfd28b67767-Paper.pdf