

Utilisation of Production Feedback for ML-Based Geometric Deviation Prediction

Jörg Brünnhäuser¹, Magdalena Dziubinska¹, Thomas Zimmermann¹, Pascal Lünemann¹, Kai Lindow¹

¹ Fraunhofer Institute for Production Systems and Design Technology, Pascalstraße 8-9, 10587 Berlin, Germany

Abstract: This paper presents a proof-of-concept approach for predicting geometric deviations of CAD designs after their production. This could help designers to recognize possible quality issues or deviations as early as possible. Several machine learning models are trained on designs and CT scanned as-manufactured parts produced by different AM processes. Results show that the used models and features have minor impact compared to the composition of the training data. The approach demonstrates the feasibility of ML-based deviation prediction to support earlier and more informed design decisions.

Keywords: geometric deviation prediction, machine learning, design feedback

1 Introduction

The transition from conventional aircraft propulsion such as jet engines to hybrid and fully electric propulsion systems motivates a fundamental shift in aircraft architecture rather than a simple incremental evolution of prior designs. Integrating electric motors, power electronics, and energy storage directly into the airframe presents unique challenges for structural design, thermal management, and system reliability. Unlike turbofan or turboprop engines that remain largely external to the primary structure of the aircraft, novel propulsion units become part of the fuselage, wing box, or empennage. This tighter integration demands a rethinking of load paths, vibration damping, and fault containment strategies.

Effective feedback-to-design mechanisms are essential to manage the increased complexity and coupling of mechanical, electrical, and thermal subsystems. By systematically capturing data during manufacturing and assembly, engineers can detect emergent issues with components and subsystems. Such feedback not only reduces development risk and cost but can also support shortened certification cycles for these next generation aircraft.

One way of feedback is the comparison of a design to its actual manufactured counterpart which could indicate the violation of tolerances or problems with either the production or the design itself. Based on such a comparison an engineer can more easily adapt the next iteration. One downside is that one must wait for the design to be produced. In our previous work (Brünnhäuser et al., 2026) we developed such a feedback to design system which processes, visualizes and eventually stores all those comparisons to a database. Based on this data we want to research whether a predictive model could be trained to give the designer a first glance at how his CAD file might look like after the production. If such a model would give a plausible prediction of deviations some problems like quality issues or the exceeding of tolerances could be solved in the current iteration instead of waiting for the next one. Therefore, it would accelerate the design process. For this reason, we intend to investigate whether an AI model can predict production deviations for specific manufacturing processes.

Initially this paper gives some examples of approaches which are predicting geometrical deviations due to the production process. After that we will present our approach how we trained a machine learning model, the details of our training data and how we conducted the training, experiments and evaluation of it. Next, we will look at the results of the experiments and the performance of the individual trained ML-models. Finally, we will discuss the results and highlight the advantages and disadvantages of our approach.

2 State of the art

Initially we conducted literature research to identify similar approaches which we could utilize for achieving our goal. There are many approaches which focus on the optimization of a single production technology, mostly additive manufacturing. In (Souza Borges Ferreira et al., 2020) an approach based on the Bayesian ELM Methodology with connectable neural networks is presented. It is able to model the deviation of additive manufacturing for a distinct AM system and it tries to adapt to different shapes or AM processes with four steps and the connectable neural networks. Another model by (Chowdhury and Anand) aims at correcting production errors with the prediction of thermal deformation. Therefore, a neural network is used based on part deformations from Stereolithography printing and after that the model is used for making corrections in the point cloud.

(Ntousia et al., 2023) created a framework which estimates the “printability of a computer-aided design (CAD) model that expresses the probability that the model is fabricated correctly via an AM technology for a specific application” (Ntousia et al., 2023). In their approach for each vertex of a mesh some features are derived and serve as the input of a neural network. This network then predicts the deviation per vertex and the printability is estimated based on those predictions and further global and part information. Another paper is also focused on the optimization of additive manufacturing with machine learning. They are using image processing to obtain the geometries of the manufactured parts. In (Li et al., 2021) conditional adversarial networks are used to improve accuracy and consistency of material extrusion. They are trained on in-situ monitoring data the CAD files itself and trained eight different models on primitive shapes. Those models can predict the deviation of freeform CAD models.

Furthermore (McGregor et al., 2022) are using a support vector regressor based on measured features to predict dimensions and qualify diverse part designs across multiple additive machines and materials. The study printed over 400 parts of three different geometries and measured five critical dimensions (like thicknesses, lengths, and diameters) on the parts which served as the label in the ML training. Based on those data and production details the support vector regressor was trained for predicting the deviations of the five dimensions and later on classify them as acceptable or not based on different strategies. (Decker and Huang, 2019) are using random forest and features which are calculated from meshes to predict geometric deviations during additive manufacturing. The random forest is trained on several features calculated for each vertex of the mesh and are about the printing bed position, curvature, surface orientation, or the linear expansion. For the training and the evaluations five simple parts were printed, laser scanned and the deviation to the CAD files were calculated. Finally, the deviation prediction was compared to the actual deviations. Another work was also based on many samples. The approach of (Baturynska and Martinsen, 2021) trained different regression models on 434 samples with a basic geometry and several dimensions were measured. A Multilayer Perceptron, decision tree regressor, gradient boosting regressor and a support vector regressor were trained and compared whether they could predict the deviations according to thickness, width, and length.

Compared to the approaches investigated we want to combine different goals which have not been researched so far. We want to focus more on the implication on the design process and leave aside the intrinsic complexity of production technologies. Our goal is to train a model which predicts geometric deviation for different manufacturing technologies. Obviously, there are many relevant parameters which could have a major impact on the outcome. For our proof of concept, we do not want to investigate those production-related parameters. Furthermore, we rely on CT-scans instead of the measurement of single dimensions or the reconstruction from pictures. Additionally, we want to use real world parts which are more complex than most of the parts used by the above approaches.

3. Methodology

In this chapter, we present the data used for training the machine learning models and outline the complete workflow, including data preprocessing, data and feature selection, and the evaluation of the proposed models. To achieve our goal of developing a model which predicts the most probable geometric deviations of a part’s design after its production, we settled on the following approach. First, data is necessary for the training of machine learning models. Therefore, some parts are selected from a project about the development of an urban air mobility electric propulsion system. Those parts are printed in several batches with different production technologies. Specifically, the parts are manufactured by three additive manufacturing processes: Fused Deposition Modelling (FDM), Stereolithography (SLA), and Selective Laser Melting (SLM). The produced parts are then scanned with computed tomography (CT) to retrieve their digital 3D models representing their actual as-manufactured geometry. We use those in combination with their corresponding CAD design files for the training of the ML-models. The combination of both makes it possible to determine and visualize geometric deviations in all three spatial directions.

In contrast to the approach proposed by (Decker and Huang, 2019) which operates on mesh representations, we use point clouds as a more suitable data representation for our use case. The scanned elements of the electric propulsion system contained an unequal number of vertices compared to the part designs, which would lead to potential issues when computing deviations between corresponding vertices. By sampling points from the mesh surfaces, the resulting dataset consists of point clouds with an equal number of points for both the scanned parts and their design counterparts. Additionally, the scans of the printed elements contained outliers, either located on the surface or appearing as isolated artifacts at larger distances from the part geometry. Working with point clouds enabled a more straightforward preprocessing pipeline, as such outliers can be efficiently removed using neighbourhood-based algorithms that identify points falling outside an allowed distance range, defined by point-to-point distances within a specified number of neighbouring points. Another important consideration concerns the registration between the CAD designs and the manufactured parts. Many registration algorithms operate by minimizing distances between corresponding points, such as variants of the Iterative Closest Point (ICP) algorithm (Pečur et al., 2025).

After obtaining the pre-processed dataset, we introduce two model training options: training on all geometries and leave-one-geometry-out training. In addition, we define multiple feature sets to compare the influence of point cloud features on

model predictions. We propose three machine learning models: XGBoost, Random Forest (RF), and XGBRF because random forest based models might be able to achieve a better result with the limited number of samples we can use for training. All details and variation of our experiments are displayed in figure 2.

3.1 Data pre-processing

The data preprocessing pipeline consists of three main steps: Poisson disk sampling of the design and manufactured geometries, outlier removal from the manufactured part point clouds, and registration of the manufactured parts to their corresponding designs.

1) Poisson Disk Sampling of Mesh Surfaces

First, the meshes are converted into point clouds. Poisson disk sampling is employed to sample the surface with 100,000 well-distributed points at approximately constant spacing (Hou et al., 2015). This number of sampling points was selected to provide adequate geometric detail for subsequent processing while remaining computationally efficient.

2) Outlier Removal from Manufactured Point Clouds

To remove outliers from the scanned manufactured parts that are not related to the part geometry itself but are introduced during the scanning process, a statistical outlier removal algorithm is applied. Points whose mean distance to their 200 nearest neighbours exceeds the global average by more than two standard deviations are removed.

3) Registration of Manufactured Parts to Designs

To enable the comparison of the point clouds and the calculation of per-point deviations between the design and manufactured parts, the parts are first aligned with each other. This is achieved by centering all parts at the origin (0, 0, 0) to obtain a common initial position, followed by registration using a global RANSAC (Random Sample Consensus) algorithm and subsequent local refinement with ICP.

The RANSAC algorithm estimates candidate point cloud transformations by randomly selecting subsets of points in each iteration and registering them to a reference point cloud. For each estimated transformation, point correspondences are established by identifying points that align within a predefined distance threshold. Points satisfying this criterion are counted as inliers. The transformation with the highest number of inliers is selected as the final alignment after either 100,000 iterations are completed or a confidence level of 0.999 is reached (Zagar et al., 2022). To further improve the registration after successful RANSAC alignment, the ICP algorithm is applied. As ICP requires a good initial alignment, it incrementally updates the transformation by minimizing squared point-to-plane distances between corresponding points. The iterations continue until the Root Mean Square Error falls below 10^{-6} or a maximum of 30 iterations is reached.

3.2 Data Partitioning for Model Training and Evaluation

In our approach to predicting deviations of the manufactured parts, we investigate the influence of the training data. The training of the different ML-models is based on three different parts from a development project of an electric propulsion system for urban air mobility. Two parts, the upper and lower part of the brackets, are responsible for the cable management and cooling on a test rig and the second one is a part of the actual electric propulsion system, the arc mounting plate. Due to the low number of only three different geometries, we settled on two different approaches to construct the training dataset.

1) All Geometries

The geometry of each of the three analysed objects is different. For each object, scans of the manufactured parts provide a dataset consisting of multiple samples associated with one of the three designs. Even for the same design, the manufactured parts are not identical, and the deviations can differ in both magnitude and direction. Due to these unique geometric characteristics, a machine learning model trained on a single geometry might not be able to make reliable predictions. In this variant, we use samples from all geometries. For the selected experimental setup, 50% of the manufactured parts are used for training and 50% for testing, while ensuring that at least one sample from each part is reserved for testing.

2) Leave-One-Geometry-Out

Typically, the performance of a machine learning model is evaluated using data that the model has not seen during training. In this experimental setting, one of the three geometries is selected as test data, while the models are trained on the remaining one or two geometries, depending on the additive manufacturing subtype used for their production.

3.3 Feature selection

The two options for model training provide a basis for analysing how machine learning model performance varies with the choice of training data. To investigate the importance of various geometrical features on the model predictions, we selected 16 distinct features:

- x, y, z – Cartesian coordinates of each point (three features).
- n_x, n_y, n_z – Surface normal vector at each point (three features).
- `distance_from_origin`: Euclidean distance of each point from the global origin (0,0,0).
- `normal_variation`: Mean standard deviation of local surface normals.
- `median_neighbor_azimuth_deg`: Median azimuth angle of 30 neighbouring points.
- `median_neighbor_elevation_deg`: Median elevation angle of 30 neighbouring points.
- `vertical_alignment`: Absolute value of the Z-component of the surface normal and representing the vertical position of a part within the production system.
- `relative_tilt_to_global`: Normalized tilt of the surface relative to the global reference frame.
- `relative_thickness`: Normalized local thickness of the geometry.
- `laplacian_norm`: Magnitude of the difference between a point and the mean position of its neighbors.
- `edge_proximity`: Relative proximity of a point to geometric edges.

Moreover, each manufactured part, and consequently all points belonging to it, is assigned one of four production detail labels corresponding to the additive manufacturing subtype used for its production: SLM, SLA, FDM without a scaling factor, or FDM with a scaling factor. To examine the influence of different sets of features, we defined seven groups and tested each combination in training of all three machine learning models. The table presents the selected feature schemes, indicating which features were removed from the full set of 16 features for the models trained under this option.

Feature scheme	Removed features
all	none
all_no_radial	distance_from_origin
geom	n_x, n_y, n_z , distance_from_origin, median_neighbor_azimuth_deg, relative_thickness
geom_with_normals	distance_from_origin, median_neighbor_azimuth_deg, relative_thickness
no_angles	distance_from_origin, median_neighbor_azimuth_deg, median_neighbor_elevation_deg, relative_thickness, n_x, n_y
no_normals	n_x, n_y, n_z , distance_from_origin
no_radial_edge_thickness	distance_from_origin, edge_proximity, relative_thickness




Figure 1: Visualization of the training and the evaluation process of our machine learning approach.

3.4 Machine Learning

To predict per-point deviations between the design and manufactured parts, three models were selected: Random Forest, XGBoost, and XGBRF, which combines gradient boosting with random forest-style tree construction. A separate model

is trained for each of the three dimensions, and the predictions are subsequently aggregated to obtain a three-dimensional per-point displacement vector (See Figure 1).

The models are trained using the selected features from the design point clouds, together with the deviations calculated for each corresponding point cloud pair along the three axes, represented as a point displacement vector. Each point cloud, and consequently all points within it, is labelled with one of the four additive manufacturing methods used for its production. To obtain predictions of the displacement vectors for a test design, the model input is defined as the features extracted from the design geometry, including the specification of the additive manufacturing process for which the predictions are made. Finally, to obtain the shape of the predicted manufactured point cloud, the estimated per-point displacement vectors are added to the design geometry.

Each of the selected machine learning models belongs to the family of tree-based algorithms. While they differ in their specific learning strategies, they share common architectural elements such as decision trees as base learners. The advantage of those models in our case is the ability to give good predictions based on a few samples. The following sections briefly describe the key characteristics of each model.

1) Random Forest

Random Forest regression constructs an ensemble of decision trees; each trained on a bootstrap sample of the training data. Within each tree, the data are recursively split at decision nodes using randomly selected subsets of features until terminal leaves are reached, where each leaf stores the mean target value of the samples assigned to it. The final prediction is obtained by averaging the outputs of all trees, which reduces variance and improves robustness compared to a single decision tree.

2) XGBoost (Chen and Guestrin, 2016)

XGBoost regression builds an ensemble of regression trees in a sequential manner. The model starts with an initial prediction, typically the mean of the training targets, and each subsequent tree is trained to predict the residual errors of the current model. The final prediction is obtained by summing the contributions of all trees, with regularization applied to control model complexity and prevent overfitting.

3) XGBRF (Divya et al., 2024 - 2024)

The XGBRF model is a hybrid model between RF and XGboost. It constructs trees using random subsampling of data and features, as in Random Forest, while using the XGBoost framework for tree construction, loss optimization, and regularization.

As the models share some architectural similarities, their hyperparameter settings can partially overlap. We used the following hyperparameters during model training. Where possible, these were chosen to be similar across all models and kept constant between experiments, in order to keep the input data and feature set selection as the main focus of evaluation in our approach (number of trees = 400, maximum depth 7, feature subsampling `max_features` = 0.8 for RF and `colsample_bytree`=0.8 for others, learning rate = 0.03 for XGBoost and 1.0 for XGBRF, subsample ratio = 0.9 for XGBoost and 0.8 for XGBRF, bootstrap sampling true for RF, tree construction method = hist for XGBRF).

3.5 Evaluation and Visualization

The evaluations of the different variations are conducted by calculating the deviations of the predictions of each trained model on the test data set to the ground truth (see figure 2). In our case the ground truth is the actual manufactured geometry with a specific production technology. For instance, we are making a prediction for the upper part of the bracket with FDM and compare it to one scan of the produced upper brackets with FDM. The comparison itself is the calculation of distance metrics between the predicted point cloud and the point cloud of the scan. We are using four different metrics: Root Mean Squared Error (RMSE), Median Absolute Distance (MAD), Chamfer Distance and Hausdorff Distance. RMSE and MAD are calculated, respectively, as the root mean square and the median of the Euclidean distances between corresponding points between the design and manufactured point clouds. Chamfer Distance and Hausdorff Distance are point cloud specific metrics that quantify geometric deviations between the predicted and reference point clouds. The Hausdorff Distance measures, in both directions, the maximum distance from any point in one point cloud to the nearest point in the other, thereby capturing the largest local deviation between the shapes. The Chamfer Distance provides a related measure, based on the average nearest-neighbour distances in both directions rather than on the maximum distance. As a result, this metric is less sensitive to individual outliers (Alibekov et al., 2024 - 2024).

Besides the numerical evaluation using quantitative metrics, a deviation representation in the form of heatmaps based on signed point-to-point distances provides a mean to visually evaluate the model prediction. To create heatmaps of deviations between the manufactured and design point clouds, point-to-point distances are computed based on closest-neighbour correspondences. Most of the actual manufactured to design deviations fall within the range of -0.9mm to 0.9mm , which is therefore used as the colour scale. The heatmap colour scale is centred at zero to represent negative (blue), zero (white),

and positive (red) deviations, and is used to compare the real deviations with those predicted by the machine learning models.

4. Results

This chapter describes the data used to train the machine-learning models, the conducted experiments, and the obtained results. As mentioned above we used three different geometries for our experiments, the two parts of the bracket and the arc mounting plate. These components were provided to us as part of the project. The parts are produced with three additive manufacturing technologies:

- Selective Laser Melting (SLM). Material: aluminum-magnesium alloy. Printer: EOS 290. Two upper part and two lower part were produced and scanned.
- Photopolymer or stereolithography (SLA) printing from PLA, printer: FORMLABS Form 4L MSLA, material: Generic Black Resin V4.1. Printing layer height 0.012 mm. One upper part and one lower part were produced and scanned.
- FDM (FDM). Printer: PRUSA MK4 enclosed, material - PLA. Layer height 0.14 mm, wall thickness 5 contours, filling density 25%. We consider it important to note that the FDM Printer was additionally calibrated to increase the accuracy of the printed part dimensions in further batches (FDM with scaling factor). Four upper and four lower parts were produced and scanned with FDM and two upper, two lower bracket parts and 7 arc mounting plates for FDM with scaling factor.

We trained the three proposed machine learning models, XGBoost, Random Forest (RF), and XGBRF, using two training data options: all geometries and leave-one-geometry-out. Furthermore, each training configuration was evaluated with different feature sets to investigate the influence of geometric point cloud features on deviation prediction which can be seen in Figure 2. The following results summarize the outcomes of these experiments under the different training and feature settings, evaluated using four selected metrics for point cloud comparison.




Figure 2. Overview over different variables of our experiments like the different set of used data, features and ML models as well as the used distance metrics to compare the results with the ground truth. The big arrows indicate the training pipeline while the small grey arrows indicate the usage of the test data set.

Table 2. Results aggregated on used ML-model and training data measured with four different metrics.

Model	Data	Chamfer [mm]	Hausdorff [mm]	RMSE [mm]	MAD [mm]
RF	leaving one geometry out	0.427	1.733	0.422	0.139
XGBoost	leaving one geometry out	0.431	1.730	0.424	0.140
XGBRF	leaving one geometry out	0.426	1.734	0.422	0.138
XGBoost	All geometries	0.348	1.472	0.148	0.123
RF	All geometries	0.348	1.475	0.148	0.123
XGBRF	All geometries	0.347	1.477	0.149	0.124

Table 2 shows the results off all trained models for each features set, the used data and the used ML model. All the different experiments are aggregated on the model type and the used training data. If we examine the impact of training a random forest, XGBoost or a XGBRFR model it is quite minor. The difference is 0.005mm Chamfer distance between the best and the least performing models for the leaving-one-geometry-out and even less for the all-geometries data set. The impact of the experiments which have seen all different geometries is much bigger. If we look at the influence of the different feature groups in table 3 we can't postulate a major difference.

Table 3. Results of the XGBoost models trained on all geometries aggregated on the different features groups. Distances measured with four different metrics.

Feature set	Chamfer [mm]	Hausdorff [mm]	RMSE [mm]	MAD [mm]
all	0.346	1.467	0.146	0.121
all no radial	0.347	1.468	0.147	0.121
coordinates only	0.352	1.493	0.152	0.126
geom	0.348	1.467	0.148	0.123
geom with normals	0.347	1.468	0.147	0.122
no angles	0.349	1.470	0.149	0.123
no normals	0.347	1.473	0.147	0.122
no radial edge thickness	0.346	1.469	0.146	0.122

Table 4. Selected results of the Random Forest model trained on geometries leaving the predicted geometry out and of the XGBoost models trained on all geometries aggregated over all feature sets and on the predicted geometry. Distances measured with four different metrics.

Model	experiment	Reference part	Production technology	Chamfer [mm]	Hausdorff [mm]	RMS E [mm]	MAD [mm]
RF	upper part	B1_registered.pcd	SLS	0.276	0.684	0.147	0.137
XGBoost	All geometries	B1_registered.pcd	SLS	0.253	0.660	0.135	0.126
RF	upper part	B1O2u_registered.pcd	FDM	0.304	0.779	0.164	0.146
XGBoost	All geometries	B1O2u_registered.pcd	FDM	0.265	0.680	0.141	0.131
RF	upper part and arc mounting plate	BG1_registered.pcd	FDM Scaling factor	0.847	4.155	1.149	0.164
XGBoost	All geometries	BG1_registered.pcd	FDM Scaling factor	0.833	4.172	0.160	0.149
RF	upper part and arc mounting plate	PhaseUnit_C_Part_2_registered.pcd	FDM Scaling factor	0.308	1.253	0.321	0.105
XGBoost	All geometries	PhaseUnit_C_Part_2_registered.pcd	FDM Scaling factor	0.274	1.269	0.113	0.096








In table 4 the results are aggregated on the prediction of a specific scanned part, model and training data set. The model predicted a geometry within the same production technology which should match a scanned part in the best case. If we look at the distances between the prediction and the ground truth, we can see a Chamfer distance around 0.3 for most parts while especially each predicted geometry for BG1 has a much higher distance to its ground truth.


We can also see in a direct comparison that the XGBoost model trained on all geometries is slightly better than the one which leaves on geometry out. To further evaluate the results, it is necessary to look at specific predictions of the best performing models in detail. At this level we can compare if the actual deviation on specific features is present in the prediction of the models. Therefore, we take a closer look at the lower part of the bracket. Specifically, we examine the BG1 examples with the highest distance to the ground truth and B2Ou with a low distance.

An overview over FDM and FDM with scaling factor is displayed in table 5. The upper two rows showing the deviations of one scanned part with FDM (first row) and FDM with scaling factor (second row) to the design. The deviations between

the individual samples in one production batch are very small. The deviation to the design is displayed colour encoded: red means a positive deviation, e.g. more material than the design, and blue for less material. The intensity of the colour scale indicates the amount of deviation—from a maximum of +9 mm and deep red to -9 mm deep blue.

Table 5. Heatmaps with positive (red) and negative (blue) deviations of the lower bracket design in comparison to an actual scan or a prediction of a ML-model.

Row	Details	Heatmap front	Heatmap back
1	Ground truth Scan of B1O2u (FDM)		
2	Ground truth Scan of BG1 (FDM with scaling factor)		
3	Prediction of lower bracket All geometries FDM features: all		
4	Prediction of lower bracket All geometries FDM (with scaling factor) features: all		
5	Prediction of lower bracket RF FDM upper part FDM features: all		

6	Prediction of lower bracket RF FDM (with scaling factor) upper part and arc mounting plate features: all		
---	---	---	---

We can now look at the specific features with a significant deviation in the ground truth. For example, we can see a slightly negative deviation on the top edge and a large positive deviation at the upper edge of the groove on the back of the FDM produced part. If we look at the corresponding edges on the backside of the predictions for the same manufacturing type in row 3 and row 5, we observe that the part trained on all geometries can predict both features while the prediction in row 5 is not. While the prediction in row 5 predicts circular deviations which aren't present in the ground truth the sample in row 3 is more accurate on the front side as well, displaying the slight red deviations on the top horizontal edge and the slight blue deviations on the vertical edges of the ribs.

If we investigate the BG1 part – where even the ground truth has a very high deviations compared to most other parts – we can see fewer deep colours on the backside. The frontside though has high deviations on the vertical edges of the ribs as well as two vertical edges on the left and right side of the part. We are missing all those details in row 6. In the prediction on row 4 we can observe the same deviations as in the ground truth but less intensive.

5. Discussion

The results indicate that both the choice of feature sets and the choice of ML models have only a minor influence on prediction quality. Across all experiments, the differences between feature schemes are within a few hundredths of a millimetre for the distance metrics. The same holds for Random Forest, XGBoost, and XGBRF: since all are tree-based methods, their predictions are very similar, with average performance differences of only about 0.005 mm between the best and worst model. For this use case, model performance is therefore dominated far more by the available data than by the specific tree-based architecture.

In the leave-one-geometry-out setting, acceptable predictions are obtained for some geometries (e.g. arc mounting plate, upper bracket), but performance drops significantly for others. This is particularly evident for the lower bracket: when this part is used only as a test geometry, and similar structural features are absent from the training set, the models fail to reproduce characteristic local deviations (e.g. at ribs or sharp edges). This supports the assumption that the models rely heavily on geometric patterns seen during training and have limited ability to extrapolate to previously unseen local features. When all geometries are used jointly for training, prediction quality improves consistently across all metrics. The aggregated Chamfer distance decreases from around 0.43 mm (leave-one-geometry-out) to approximately 0.35 mm, and RMSE and MAD also drop. Visual inspections of deviation heatmaps confirm this trend: models trained on all geometries reproduce typical deviation patterns—such as slightly negative deviations at top edges or positive deviations in groove regions—more accurately than models trained with one geometry excluded. At the same time, this setup increases the risk of overfitting to the limited set of parts. Given the small number of distinct geometries, it is not possible to clearly distinguish between genuine generalisation of geometry-specific deviation patterns. A more reliable assessment of generalisation would require a larger and more diverse dataset.

In cases where the ground-truth deviations are unusually large (e.g. part BG1 and two other parts), prediction errors also increase. This may indicate that the models struggle with rare or extreme deviation patterns that are underrepresented in the training data. However, the results indicate that, even with a limited set of features directly linked to point-wise changes between design and manufactured parts of one part type, the models can predict deviations for unseen parts with higher accuracy—provided that the manufacturing process produces similar deviation patterns for all parts. This can be expected as a potential limitation of the machine learning tree-based models, and in the future the use case can be examined with deep-learning models, which use the local and global geometry-based information in the training process.

6. Conclusion

This work presented a proof-of-concept approach for predicting geometric deviations between CAD designs and as-manufactured parts using CT-based point clouds and tree-based machine learning models. The approach demonstrates that it is possible to make reliable predictions of upcoming deviations for specific AM technologies, providing a basis for

designers to adapt their designs and avoid problems that might arise later in the production process. The results show that Random Forest, XGBoost and XGBRF achieve similar accuracy and can qualitatively reproduce typical deviation patterns for different AM processes. The dominant performance factor is the composition of the training data rather than the specific model or feature set: models trained on all geometries clearly outperform leave-one-geometry-out setups, but may overfit to the limited set of parts.

The study highlights both the potential and the current limitations of ML-based deviation prediction for design feedback. To strengthen the contribution, we want to adapt a more principled training-data selection strategy which we couldn't apply here because the lack of training data. Rather than selecting training data primarily at the part level (e.g. all geometries vs. leave-one-geometry-out), we will adopt a structure-aware approach that explicitly accounts for shared geometric characteristics and typical deviation patterns if training data allows it. To achieve robust generalisation to new parts and processes, future work must extend the dataset to more diverse geometries, incorporate process parameters, and explore more expressive point-cloud-based learning architectures. In the future we want to train and test different ML-models like support vector regression, feed forward neural network or use deep learning models specialized on point clouds. One complexity which wasn't addressed in this paper is the consideration of different parameters for each AM technology or even each 3D printer. The impact might be very big, so this should be taken into account in future work.

Funding: This research was funded by Federal Ministry for Economic Affairs and Climate Protection based on a resolution passed by the German Bundestag, grant number 20L2108B

ReferencesReferences

- Alibekov, U., Staderini, V., Ramachandran, G., Schneider, P., Antensteiner, D., 2024 - 2024. Evaluation of 3D Point Cloud Distances: A Comparative Study in Multi-Point Cloud Fusion Environments, in: Proceedings of the 19th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications. 19th International Conference on Computer Vision Theory and Applications, Rome, Italy. 27.02.2024 - 29.02.2024. SCITEPRESS - Science and Technology Publications, pp. 59–71.
- Baturynska, I., Martinsen, K., 2021. Prediction of geometry deviations in additive manufactured parts: comparison of linear regression with machine learning algorithms. *J Intell Manuf* 32, 179–200. <https://doi.org/10.1007/s10845-020-01567-0>.
- Brünnhäußer, J., Dziubinska, M., Zakheer, U., Bilous, V., Zimmermann, T., Joost, R., Lindow, K., 2026. Enhancing Product Design in Electric Aviation through Digital Twins and Production Feedback Integration. [Manuscript submitted and accepted for publication]. *Engineering Proceedings Journal - EASN 2025 proceedings*.
- Chen, T., Guestrin, C., 2016. XGBoost: A Scalable Tree Boosting System. <https://doi.org/10.48550/arXiv.1603.02754>.
- Chowdhury, S., Anand, S. Artificial Neural Network Based Geometric Compensation for Thermal Deformation in Additive Manufacturing Processes, in: Volume 3: Joint MSEC-NAMRC Symposia 2016.
- Decker, N., Huang, Q., 2019. Geometric Accuracy Prediction for Additive Manufacturing Through Machine Learning of Triangular Mesh Data. <https://doi.org/10.1115/MSEC2019-3050>.
- Divya, S.V., Venkadesh, P., Pavithra, R.J., M.R., A., G, L., A, L., 2024 - 2024. An Integrated XGBRF Framework for Early Identification of Parkinson's Disease, in: 2024 International Conference on Emerging Research in Computational Science (ICERCS). 2024 International Conference on Emerging Research in Computational Science (ICERCS), Coimbatore, India. 12.12.2024 - 14.12.2024. IEEE, pp. 1–9.
- Hou, W., Xu, Z., Qin, N., Xiong, D., Ding, M., 2015. Surface reconstruction through poisson disk sampling. *PloS one* 10, e0120151. <https://doi.org/10.1371/journal.pone.0120151>.
- Li, L., McGuan, R., Isaac, R., Kavehpour, P., Candler, R., 2021. Improving precision of material extrusion 3D printing by in-situ monitoring & predicting 3D geometric deviation using conditional adversarial networks. *Additive Manufacturing* 38, 101695. <https://doi.org/10.1016/j.addma.2020.101695>.
- McGregor, D.J., Bimrose, M.V., Shao, C., Tawfick, S., King, W.P., 2022. Using machine learning to predict dimensions and qualify diverse part designs across multiple additive machines and materials. *Additive Manufacturing* 55, 102848. <https://doi.org/10.1016/j.addma.2022.102848>.
- Ntousia, M., Fudos, I., Moschopoulos, S., Stamati, V., 2023. Predicting geometric errors and failures in additive manufacturing. *RPJ* 29, 1843–1861. <https://doi.org/10.1108/RPJ-11-2022-0402>.
- Pečur, T., Bosché, F., Cerniauskas, G., Mill, F., Sherlock, A., Yu, N., 2025. Prototype pipeline modelling using interval scanning point clouds. *Adv. Manuf.* 13, 444–461. <https://doi.org/10.1007/s40436-024-00515-y>.
- Souza Borges Ferreira, R. de, Sabbaghi, A., Huang, Q., 2020. Automated Geometric Shape Deviation Modeling for Additive Manufacturing Systems via Bayesian Neural Networks. *IEEE Trans. Automat. Sci. Eng.* 17, 584–598. <https://doi.org/10.1109/TASE.2019.2936821>.
- Zagar, B.L., Yurtsever, E., Peters, A., Knoll, A.C., 2022. Point Cloud Registration With Object-Centric Alignment. *IEEE Access* 10, 76586–76595. <https://doi.org/10.1109/access.2022.3191352>.