

HIGH-RESOLUTION SHAPE DEFORMATION PREDICTION IN ADDITIVE MANUFACTURING USING 3D CNN

Benjamin Standfield
Denis Gračanin

Rongxuan Wang
Zhenyu Kong

Department of Computer Science
Virginia Tech
2202 Kraft Drive
Blacksburg, VA 24060, USA

Department of Industrial and Systems Engineering
Virginia Tech
1145 Perry Street
Blacksburg, VA 24061, USA

ABSTRACT

Additive manufacturing (AM) processes usually have lower geometric quality and when compared to subtractive manufacturing processes. However, AM processes are seeing more use in the industry because they are both affordable and flexible. To address the lower geometric quality and reduced reliability drawbacks, 3D Convolutional Neural Networks (CNN) were developed and used to predict the deformations from the ideal sliced 3D object. The developed 3D CNN were tested on a live dataset consisting of 50 3D printed, 3D scanned, and aligned objects. The linear spatial resolution of these predictions is improved to $150\mu\text{m}$ with a sampling frequency of 166 units per inch compared to the standard peak resolution of 64 units across an axis. Results indicate that using the described approach provides better predictors of part geometry than the original STereoLithography (STL) file defining the part. An average increase of the F1 measure is 0.0644 over using the STL.

1 INTRODUCTION

Since its initial development in 1986 (Hull 1986), additive manufacturing (AM) processes have been inferior to the more established subtractive manufacturing (SM) processes in terms of quality, and reliability (Zhang et al. 2018).

On the one hand, AM provides a more flexible manufacturing pipeline allowing the creation of shapes and geometries that were either difficult or impossible to create with traditional SM methods thereby making it an attractive option for industry. Because of the benefit of AM processes over SM processes in the industry, it is worth attempting to mitigate the drawbacks of AM processes by increasing accuracy and, thereby, quality and reliability of AM-produced parts.

On the other hand, hobbyists have also taken a liking to 3D printing technologies as some of these technologies are affordable and support rapid prototyping and visualization of 3D structures. Hobbyists may also be more productive and accurate as increased accuracy of AM processes reduces the chance of wasted material and rework.

Only recently have AM processes found their way into the industry. Quality is still a key issue inhibiting the realization of AM processes in areas where there are low geometric tolerances are necessary, such as high-performance injection molds, aerospace applications, and optical structures (Gao et al. 2020; Shen et al. 2019; Pfuhl and Degünther 2021). AM processes and equipment have become simple and cheap enough for the hobbyist. However, cheaper 3D printers tend to have larger error than their commercial and industry-approved counterparts (Duan et al. 2018). Additionally, trends towards social and cloud computing lead to cloud factories where the client's needs and quality requirements drive production (Shen et al. 2019). In order to keep up with these trends, AM processes need to increase production quality.

The capability of a prediction model to determine what and where the geometric defects will occur results in saving time and money. These savings are very important to industry acceptance as a standard manufacturing process and hobbyists who use AM to create physical models such as sculptures. 3D Convolutional Neural Networks (CNN) models were developed and used to predict the deformations from the ideal sliced 3D object.

To develop and test this CNN model, a live data set was created by 3D printing and scanning a set of objects to serve as input and output data. After printing and scanning, the scanned object will be aligned and discretized to be used as input and output for the CNN model. An algorithm external to the CNN is responsible for logically breaking up smaller pieces of the larger 3D space so that higher resolution data can be passed into the model. That same algorithm then stitches together the smaller resulting spaces and computes performance metrics.

The main contribution of the proposed approach is increased spatial resolution of 3D CNN model-based predictions. Specifically, the contributions are as follows:

- A $150\mu\text{m}$ resolution prediction model is constructed using a 3D CNN which is, to the best of our knowledge, the highest resolution 3D CNN-based method for error prediction in the field. The resolution of the method and model can be changed to any desired resolution so long as the resolution of the data allows.
- A new methodology for 3D CNN-based prediction is proposed that runs on subsections of the 3D space and stitches the results of multiple predictions together. This methodology reduces the GPU memory strain of predicting with high spatial resolution as it is no longer necessary to load the entire 3D space into GPU memory. The methodology is tested on a live data set of crescent shapes and cylinders on top of a plate, which is then scanned with a 3D scanner to be used as output data while the original STereoLithography (STL) file is used as input data.
- The new methodology shifts the problem from 3D spaces not fitting into memory to another issue, the resolution of the model per unit of space versus the GPU memory. In the case of high model resolution, the input space to the model needs to be large to accurately predict the output space.
- From a practical point of view, one can expect less wasted material and time due to reprinting or subtractive post-processing stages in the manufacturing pipeline.

The remainder of the paper is organized as follows. Section 2 outlines work done in the field of error prediction and 3D scanning. Section 3 presents the proposed methodology for addressing the research problem. Section 4 documents the process of generating the dataset used to address the research project. Section 5 shows the results of the experimental procedure presented in the research methodology. Finally, the paper is concluded, and the future work as well as limitations, is presented in Section 6.

2 RELATED WORK

To address the lack of geometric accuracy, different approaches have been employed to model the error using deep learning (Shen et al. 2019) and statistical models (Wang et al. 2017). Customized models of deep learning models, such as PredNet (Shen et al. 2019), have been developed for prediction which has been shown to be effective on 2D simulated data. This work was then extended to predict a $64 \times 64 \times 64$ space with similar accuracy (Zhao et al. 2019). Additionally, Multi-layer Perception models and traditional CNN's have been employed (Baturynska, Semeniuta, and Wang 2019) to predict geometric features such as length as opposed to point cloud data. Furthermore, Random Forest models have been employed to predict point cloud data represented as polar coordinates (de Souza Borges Ferreira et al. 2020).

Statistics-based models have been developed and used to model an array of problems in various fields. Within AM error prediction, various models have been used to model geometric characteristics of AM processes. Weighted optimization models have been employed (Chowdhury et al. 2018) to predict point cloud data. Finite element models have also been employed in predicting wire arc AM processes (Casuso

et al. 2021). A new model was also developed using polar coordinates to represent points on an AM process, however, the method does not cover Z direction deformation (Huang et al. 2015). The methodology was later modified in order to account for this Z axis deformation (Jin et al. 2015). A skin model-based modal decomposition approach was also employed to classify errors and predict geometric characteristics (Huang et al. 2018).

3D scanning, on the other hand, is a powerful tool to inspect the quality of an AM part by converting physical geometry to point cloud data. Other than the geometry, it can also reveal the surface texture and defects of a print. In general, non-contact scanning methods are preferred since they are non-destructive. These methods include laser triangulation, photogrammetry, and structured light scanning (SLS) (Dorsch et al. 1994; Mikhail et al. 2001; Geng 2011).

Wang et al. (2021) analyzed and explained the principle of each method in detail and concluded that SLS is the most suitable technology for 3D printing-related applications. They developed a customized 3D scanner that yields high-spatial resolution data with a high-speed and use it to inspect the metal 3D printed part's surface. Chenang et al. (2021) used the scanned result from SLS to train a deep learning model to predict the surface roughness based on image data. Law et al. (2022) used SLS to inspect the surface geometry of bio 3D printing. A previous work used the same method to scan the 3D printed clay object and analyze the distortion (Wi et al. 2020).

3 METHODOLOGY

Our goal is to determine whether deep learning can be employed for high-resolution error prediction by breaking a high-sampled 3D mesh into smaller samples, performing independent predictions, and stitching the results back together. Breaking the meshes into smaller pieces has two key advantages:

1. An increased number of training samples generated from a single object, making it easier for a model to converge.
2. An increased resilience to new unseen shapes as the models would predict a small section of the 3D object based on the surrounding geometry.

However, computation time and data complexity are also increased for both training and prediction as multiple GPU passes are necessary for a single option, unlike other methodologies.

To validate the methodology, two CNN models are employed from the literature that has had high success in lower resolution predictions on voxelized data, namely PredNet (Shen et al. 2019) and the 3D CNN (Zhao et al. 2019). Additionally, a model is developed based on the Residual Neural Network (ResNet) architecture (He et al. 2015) for predictions on the subspace data. Each of these models is adapted to support the new input and output data structures required for supporting the sub-space method. An overview of the methodology is summarized in Figure 1.

Any method attempting to model error on a process needs to consider sources of errors. In the case of a geometric error for additive Fused Deposition Modeling (FDM) processes, it has been shown that there are six key sources of errors (Decker and Huang 2019) — X position, Y position, Z position, surface orientation, curvature, and linear thermal expansion effects. A 3D CNN trained on the 3D sub-space voxel data will not have access to the global X , Y , and Z positions because the 3D segments do not represent the entire space unlike the conditions present in previous literature.

The model may be capable of learning the relationship between the input and output space regarding the surface orientation, curvature, and linear thermal expansion effects as the former two can be determined by surrounding geometry and the latter can. To address the lack of global spatial coordinates, the input data to each model should be modified to include the X , Y , and Z positions, each of which should have a larger impact on the geometric error as the positions deviate more from the center of the print bed, in addition to the 3D voxel data.

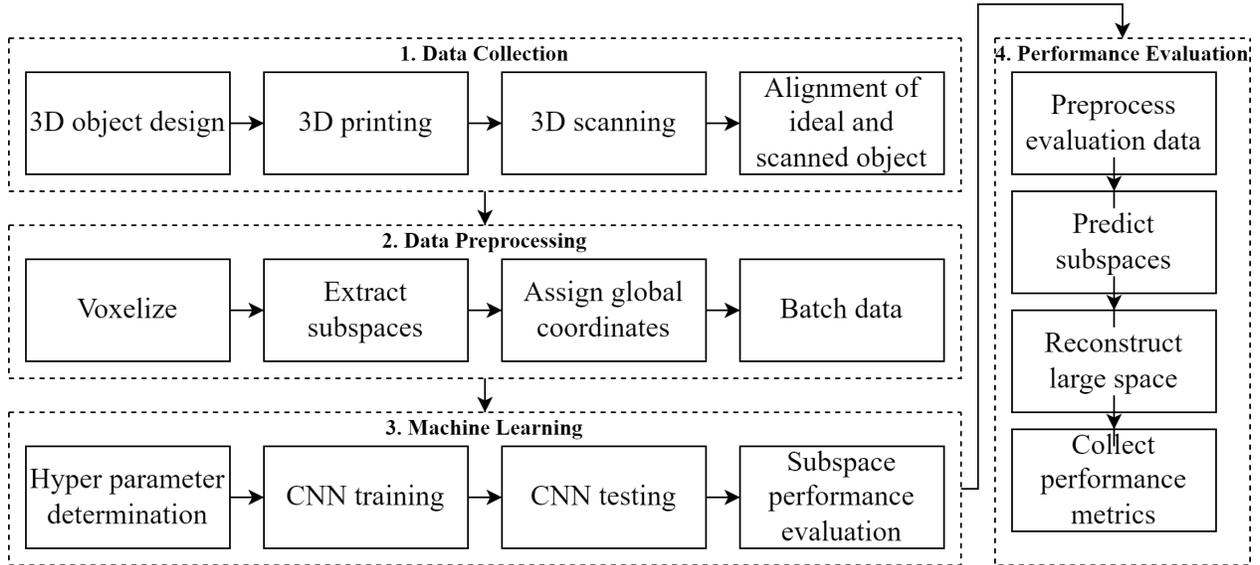


Figure 1: Methodology overview — After data collection and preprocessing, machine learning is applied to train CNN models that are then evaluated.

To validate the methodology of this paper, an experiment was conducted using three deep learning models, one adapted from Shen et al. (2019)’s work, one adapted from Zhang et al. (2018)’s work, and a new 3D CNN model based on the ResNet architecture. Each of these models was trained, tested, and validated using the training data from the scanned objects using cross validation. Eight of ten printed objects were used for training, one printed object was used for testing and validation, and one printed object was used for quantifying the performance of the model. During the testing and validation step, a threshold was identified to binarize the results of the deep learning model’s prediction by identifying the mean and standard deviation of all voxels that should have been marked as solid and all voxels that should have been marked air. After identification, a midpoint between these two means was selected by incrementing or decrementing each mean by its standard deviation and then selecting the mid point when the two values overlap. During the quantification of performance phase, the model was run on the remaining object to generate a 3D object prediction and determine F1 score.

The PredNet model (Shen et al. 2019) is a $32 \times 32 \times 32$ resolution network designed to predict 3D spaces. The revised architecture base don this network is shown in Figure 2. The original network utilizes a series of convolution layers followed by max pool layers on the first half of the network. The second half of the network concatenates the results of the first half and a deconvolution layer. To adapt the network to fit the proposed data format, six changes are made to the original architecture:

1. The second half of the network utilizing deconvolutions and concatenation operations is removed.
2. The second branch of layers is added to process the spatial data.
3. A flatten operation is added after the last pool and is then combined with the output of the spatial data branch into a dense layer with a Rectified Linear Unit (ReLU) activation function.
4. A final dense layer is added with a sigmoid activation function which is then fed into a reshape layer to resize the output dimensions to $3 \times 3 \times 3$.
5. Sigmoid focal cross entropy is used as opposed to the original improved cross entropy loss function.
6. Input size of $25 \times 25 \times 25$ with a batch size of 192 is used instead of the intended $32 \times 32 \times 32$ with a batch size of 64. Additionally, an input size of $35 \times 35 \times 35$ with a batch size of 64 is used to observe the impact of batch size on the model performance.

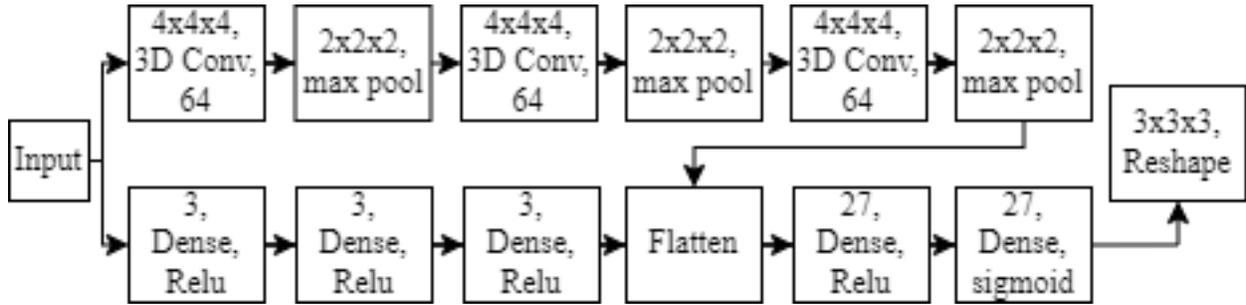


Figure 2: The proposed model architecture based on PredNet (Shen et al. 2019).

Because the CNN developed by Zhao et al. (2019) is also based on Shen et al. (2019)’s work, the modified network based on Zhao et al.’s architecture (Figure 3) is similar to Shen et al.’s modified architecture. Unlike Shen et al.’s PredNet, the network intends to take a much higher resolution space, $64 \times 64 \times 64$, at the cost of the batch size being much lower at ten. Six changes were also made to Zhao et al.’s original architecture to accommodate the methodology of this paper.

1. The second half of the network utilizing deconvolutions and concatenation operations is removed.
2. The second branch of layers is added to process the spatial data.
3. A flatten operation is added after the last pool and is then combined with the output of the spatial data branch into a dense layer with a ReLU activation function.
4. A final dense layer is added with a sigmoid activation function which is then fed into a reshape layer to resize the output dimensions to $3 \times 3 \times 3$.
5. Sigmoid focal cross entropy is used as opposed to the original improved cross entropy loss function.
6. Input size of $63 \times 63 \times 63$ with a batch size of ten is used instead of the intended $64 \times 64 \times 64$ with a batch size of ten.

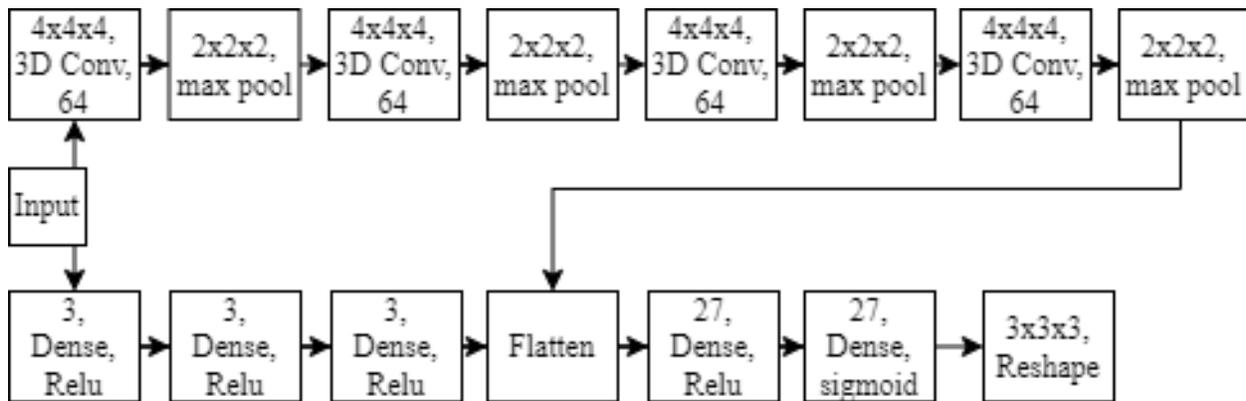


Figure 3: The proposed model architecture based on a 3D CNN model (Zhao et al. 2019).

The model proposed by this paper is based loosely on the ResNet (He et al. 2015) architecture (Figure 4). The 2D convolutions are replaced with 3D convolutions while keeping the number of channels after each ResNet Convolution operation small. The model takes sub-spaces of size $35 \times 35 \times 35$ with a batch size of 64 to output a $3 \times 3 \times 3$ subspace. Mean Squared Error (MSE) was selected as the loss function as it outperformed the sigmoid cross entropy loss function in prior experiments when tested using an array of different α and γ values. Experiments also showed that high numbers of filters in deeper layers had

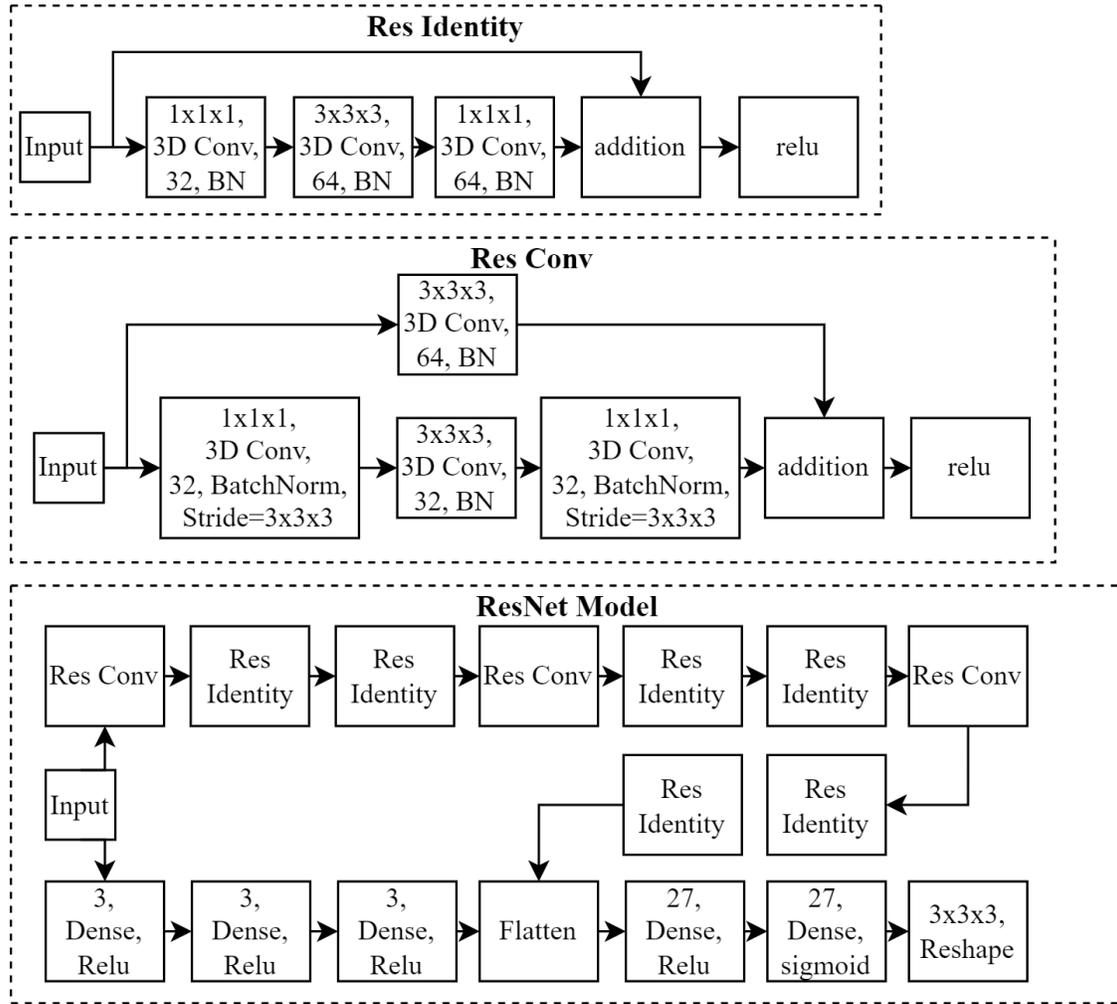


Figure 4: The proposed model architecture based on ResNet (He et al. 2015).

a negative effect on model accuracy. To address this, filters were kept relatively constant throughout the model.

4 DATASET GENERATION

To generate the dataset, 50 3D models are first developed using Solidworks, a 3D CAD program. After development, they are printed by first slicing using Cura software, then 3D printed on the Monoprice Mini V2 3D printer. Next, they are scanned using a Metron 3D scanner and a turn table to register ten 3D scans over 360 degrees into a single object. After scanning, the scanned object and ideal STL file are aligned with Cloud Compare software using an iterative closest point algorithm on the STL file and the scanned object. Finally, ten randomly selected STL files and scanned object files are sampled at 200 points per inch to generate the voxels, and 3D cubes are cut out from the input and output objects in order to serve as the training, testing, and validation data.

Each of the training objects took the same general shapes as depicted in Figure 6a, a crescent of varying shape resting on a plate along-side a cylinder. The plate measures 1" × 1" × 0.25", and the cylinder and crescent both measure 0.25" tall. The purpose of the cylinder is to help align the scans during the scanning phase. Otherwise, the scanning software may have difficulty aligning multiple scans when the crescent is

either non-existent or a full circle. The crescent shape resting in the middle of the plate is two extruded overlapping circles with one subtracted from the other. Variation is introduced by having one of the cylinders move along the X -axis before being extruded and subtracted. After development, each object was printed, measured to obtain an approximation of error on each axis, and scanned to be used in the dataset.

A customized SLS scanner was built and used for scanning the sample after the print using the HP5 Scan software. The scanner was calibrated to yield a $65.3\mu\text{m}/\text{pixel}$ spatial resolution. Before the scan, the parts are applied to a uniform anti-reflection spray (ARDROX 9D1B developer), as shown in Figure 6b, then positioned on a rotational table, which is automated with the scanner (Figure 5). The final geometry is a fused result of 10 scans, each from a different angle (36° apart).



Figure 5: Experimental setup of the scanning process. A 360 degree turntable is used along with a Metron projector and HP 3D HD camera pro to facilitate the scanning process.

The result of the scanning process is shown in Figure 6c. The input and output objects needed to be aligned so the input 3D space and output 3D space from the perspective of the CNN would represent the same position with respect to each object's origin. To accomplish this, the 3D printed object was aligned with the STL file using a four-point iterative closest point algorithm in the Cloud Compare software. After alignment, the two objects are saved as `.obj` files in pairs for training data collection in Blender.

The four points chosen were the four bottom corners of the input and output objects for two reasons. First, the first layer of the 3D print will have the least error due to the flat bed the print takes place on and the lack of error from shrinkage in the Z axis. Second, the corners of the mesh are easy to identify and select in both the input STL and the scanned mesh.

Pre-processing of data took place in two stages, the binarization of data objects and the extraction of 3D samples and batching of samples. Binarization took place sequentially for each 3D object and began by loading the first loading a 3D object into Blender. Each axis was sampled 200 times over 1.25 inches to leave sufficient empty space surrounding the object.

For each point within the sampled 3D space, a line was drawn in the opposite direction of the origin to determine whether the point was inside or outside the mesh using the dot product of the normal of the first face collided. If there was no collision, the point was deemed to be outside of the object. If the point is the origin, and hence the direction does not exist, the point was deemed to be inside the object. The result of this operation is a 3D array of indicator variables shown in Figure 6d. However, the X , Y , and Z positions of the respective point were appended to each indicator variable for future training.

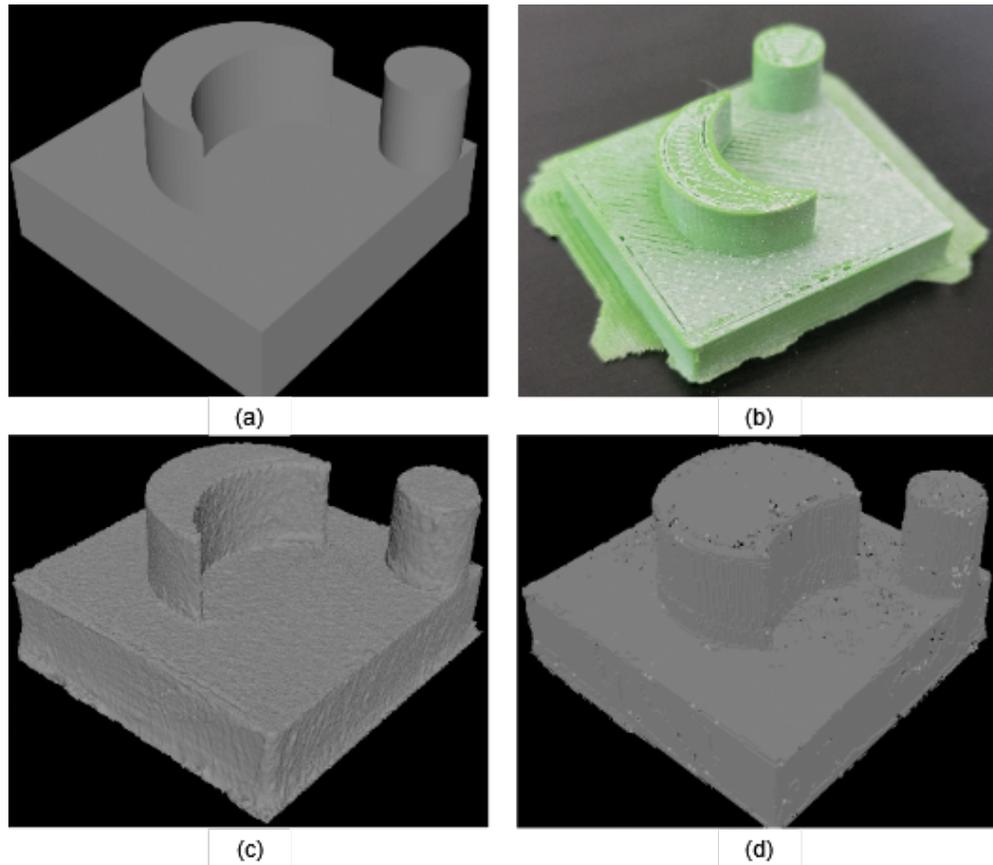


Figure 6: (a) STL file. (b) Printed and painted 3D object. (c) Scanned 3D object. (d) Voxelized 3D object.

Two additional filters were added to reduce the size of the dataset and increase the quality of the training data. The first filter removed the sample when the amount of filled space represented by the 3D sub space is below 10% or above 90%. This reduces the overfitting due to the large amounts of space within and outside the model thereby, increasing the quality of the predictions. The second filter removes 75% of samples with the only goal of reducing the dataset size. The final result of this operation was a set of 4D arrays, one for each 3D object, saved as binary files.

The last preprocessing step, utilized the saved binary files from the previous step to create the dataset files. Using the same sampling frequency, 4D arrays are extracted from both the input and output objects using the same center coordinate but a different radius of extraction. A large 3D array is extracted from the input space by leaving off the position variables. The large 3D array is then flattened into a 1D array, and the three position variables of the center coordinate are then appended to the array.

As for the output, a small 3D array is extracted from the output space with the same center position as the large input array by leaving off the indicator variables. This results in a 1D array input and a 3D array output for training. This process is repeated for each position in the 4D input array and for each pair of input and output objects. Positions along the edge of the 4D space where the extracted input 3D space would go outside the bounds of the 4D array are omitted from the dataset. The final result of this pipeline is 910 batches of 64 data samples totaling 58,240 samples.

5 RESULTS

The models based on Shen et al.'s and Zhao et al.'s architectures both used sigmoid focal cross entropy loss functions (Lin et al. 2017). This loss function has been shown to work well for increasing model accuracy

in the case of unbalanced data. There is a filter in the data generation phase to reduce the impact of data imbalance. However, the sigmoid focal cross entropy was still used to stay true to the original models' architectures. The sigmoid focal cross entropy loss function take two parameters, α and γ , to determine how much weight to assign to each class.

To determine the optimal parameters for each network, a grid search was employed over an array of potential values. To find α , a number between 0 and 1, 20 equally spaced values between 0 and 1 were tested. To find γ , 10 equally spaced values between 0 and 5 were tested. The optimal combinations were determined by evaluating each model to find the one with the highest F1 in the testing and validation phase.

A summary of the results is provided in Table 1. After testing each of the combinations, an α of 0.65 and a γ of 2.5 had the highest F1 measure for the Shen-based 35³ input size model with an F1 measure of 0.9283. The Shen-based 25³ input size model had a peak F1 of 0.8987 using an α of 0.45 and a γ of 4. Lastly, the optimal α and γ for the mode based on Zhao et al.'s architecture was determined to be 0.65 and 1 respectively, and the resulting F1 measure is 0.3312.

Table 1: Optimal α and γ for the adapted architectures.

Model	batch size	input size	α	γ
Shen et. al-based	64	35	0.65	25
Shen et. al-based	192	25	0.45	4
Zhao et. al-based	10	63	0.65	1

After printing 50 objects, each was measured on the X, Y, and Z axes, and the middle of the crescent to determine how much their lengths deviate from the targeted dimensions. The summary of these measurements can be found in Table 2. The largest error came from the Y axis with an average difference of 0.02483 from the target length. Given that the model resolution is approximately $\frac{1}{166}$ of an inch, this equates to four voxels in the Y dimension. The smallest error was present in the crescent and less than one voxel represents the error in that dimension. Overall, there is a significant error from the models' perspective and the voxelized should be capable of reflecting it.

Table 2: Summary of error for each dimension of the data (inches).

Statistic	x	y	z	crescent
Ideal	1	1	0.5	0
Average	1.006468	0.97517	0.48063	-0.00086
Average Δ	0.006468	-0.02483	-0.01937	-
Standard Deviation	0.0129	0.0052	0.0036	0.0299

$$F_1 = \frac{TP}{TP + \frac{1}{2}FP + FN} \quad Precision = \frac{TP}{TP + FP} \quad Recall = \frac{TP}{TP + FN} \quad (1)$$

F1 was calculated using the scanned object as the ground truth. A true positive signifies that the scanned result and prediction both indicate an area is solid. True negative signifies that both objects have an area marked as air. A false positive indicates that the scanned result is air, but the prediction is solid. A false negative indicates that the ground truth is solid while the prediction is air. Precision, on the other hand, is a representation of how frequently the model generates false positive results. A high precision value indicates a lower amount of false positives. Lastly, recall, is an indicator of the amount of false negatives made by the predictor. A high recall indicates a low amount of false negatives. F1, precision, and recall each take a value between 0 and 1 with 1 being ideal and the equations are provided in equation 1.

The results of training and testing each model are summarized in Table 3. The architecture based on Shen et al. (2019), ResNet, and Zhao et al. (2019). are accompanied by the ideal case for performance comparison. The ideal set of metrics consists of the performance of using the ideal STL file as an indicator

for the result of the 3D print operation. In general, the Shen et al. (2019) architectures out-perform the remaining models and the STL file comparison. The ResNet-based architecture comes close to the performance of the Shen et al.-based network in most metrics and pulls ahead in the category of recall. The network based on Zhao et al. has the lowest performance out of the models of comparison. Additionally, due to its batch size of ten, it also takes the longest to predict a complete object. A visualization of the results can be found in Figure 7. The Architecture based on Zhao et al. is omitted as the model assumes the entire 3D space is solid, resulting in a large box.

While the Shen et al. and ResNet-based network take about three minutes to predict a 200^3 object, the Zhao et al. takes approximately two hours, making it by far the slowest of the three models. The ideal object serves as a baseline to quantify improvement. Using the ideal as a comparison, it can be said that Shen et al. and ResNet models both outperform using the ideal file as an indicator for printed geometry with a spatial resolution of approximately $150\mu\text{m}$.

Table 3: Model comparison.

Model	F1	Precision	Recall	Accuracy
shen64	0.9324	0.9110	0.9548	0.9656
shen192	0.9285	0.9197	0.9374	0.9692
resnet64	0.9254	0.9057	0.9460	0.9621
resnet192	0.9099	0.8454	0.9851	0.9584
zhao10	0.1436	0.0915	0.3337	0.6387
ideal	0.8597	0.7822	0.9542	0.3889

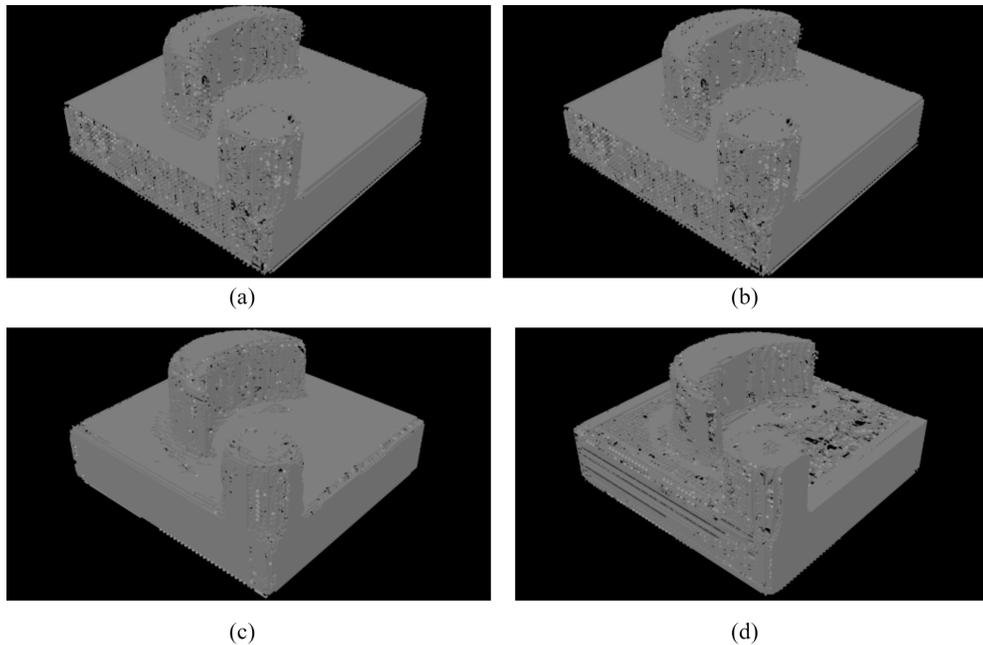


Figure 7: Prediction results: (a) Shen batch size 64. (b) Shen batch size 192. (c) ResNet batch size 64. (d) ResNet batch size 192.

6 CONCLUSION

We demonstrated that it is possible to perform predictions at much higher resolutions than what was previously done before (a resolution of 64^3). Additionally, we used a live data set as opposed to the simulated ones used in the literature. With more than tripled linear spatial resolution methodology for error prediction, one can expect a much high degree of precision when estimating part geometries. Based on the results of the aforementioned methodology, one can expect accurate predictions far surpassing that of using the ideal STL as an estimation of the printing process result. Furthermore, the methodology is shown to be effective with spatial resolutions as low as $150\mu\text{m}$.

Benefits aside, there is a new set of problems brought about by this new methodology. The key problem for high-resolution predictions before was the lack of memory available for loading large 3D arrays. More specifically, prediction resolution is scaled directly with spatial resolution. With this new methodology, the higher the spatial resolution used for predictions, the more neighboring pixels are required for an accurate prediction.

Because error prediction in 3D printing is typically a precursor to error compensation, in the future, we would like to build a high spatial resolution error compensation method following a similar methodology. Using the spatial coordinates within each of the training samples, it should be feasible to create a compensation model with similar architectures. In addition to error compensation, extended reality (XR) can be used to build a real-time estimator for the 3D printing process. In conjunction with XR, which is considered to be another rapid prototyping process, a tool can be developed and user-tested in hopes of reducing the amount of time spent in the modeling-printing loop before the 3D object conforms to the desired tolerance.

REFERENCES

- Baturynska, I., O. Semeniuta, and K. Wang. 2019. "Application of Machine Learning Methods to Improve Dimensional Accuracy In Additive Manufacturing". In *Advanced Manufacturing and Automation VIII* (1st ed.), edited by K. Wang, Y. Wang, J. Strandhagen, and T. Yu. New York, New York, USA: Springer.
- Casuso, M., F. Veiga, A. Suárez, T. Bhujangrao, E. Aldalur, T. Artaza, J. Amondarain, and A. Lamikiz. 2021. "Model for the Prediction of Deformations in the Manufacture of Thin-Walled Parts by Wire Arc Additive Manufacturing Technology". *Metals* 11(5):678–694.
- Chenang, L., W. Rongxuan, Z. Kong, B. Suresh, J. Chase, and F. James. 2021. "Real-Time 3D Surface Measurement in Additive Manufacturing Using Deep Learning". In *2019 International Solid Freeform Fabrication Symposium*. August 12th-14th, Austin, Texas, USA, 225-237.
- Chowdhury, S., K. Mhapsekar, and S. Anand. 2018. "Part Build Orientation Optimization and Neural Network-Based Geometry Compensation for Additive Manufacturing Process". *Journal of Manufacturing Science and Engineering, Transactions of the ASME* 140(3):1–15.
- de Souza Borges Ferreira, R., A. Sabbaghi, and Q. Huang. 2020. "Automated Geometric Shape Deviation Modeling for Additive Manufacturing Systems via Bayesian Neural Networks". *IEEE Transactions on Automation Science and Engineering* 17(2):584–598.
- Decker, N., and Q. Huang. 2019. "Geometric Accuracy Prediction for Additive Manufacturing Through Machine Learning of Triangular Mesh Data". In *ASME 2019 14th International Manufacturing Science and Engineering Conference*. June 10th-14th, Austin, Texas, USA, 1–9.
- Dorsch, R. G., G. Häusler, and J. M. Herrmann. 1994. "Laser Triangulation: Fundamental Uncertainty in Distance Measurement". *Applied Optics* 33(7):1306–1314.
- Duan, M., D. Yoon, and C. E. Okwudire. 2018. "A Limited-Preview Filtered B-Spline Approach to Tracking Control – With Application to Vibration-Induced Error Compensation of a 3D Printer". *Mechatronics* 56(1):287–296.
- Gao, Z., D. Sossou, and Y. Zhao. 2020. "Machine Learning Aided Design and Optimization of Conformal Porous structures". In *Proceedings of the ASME 2020 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*. August 17th-19th, Virtual, Online, 1–9.
- Geng, J. 2011. "Structured-Light 3D Surface Imaging: a Tutorial". *Advances in Optics and Photonics* 3(2):128–160.
- He, K., X. Zhang, S. Ren, and J. Sun. 2015. "Deep Residual Learning for Image Recognition". Technical Report abs/1512.03385, Computing Research Repository, Cornell University, Ithaca, New York, USA.
- Huang, Q., J. Zhang, A. Sabbaghi, and T. Dasgupta. 2015. "Optimal Offline Compensation of Shape Shrinkage for Three-Dimensional Printing Processes". *IIE Transactions* 47(5):431–441.

- Huang, Z., J.-Y. Dantan, A. Etienne, M. Rivette, and N. Bonnet. 2018. "Geometrical Deviation Identification and Prediction Method for Additive Manufacturing". *Rapid Prototyping Journal* 24(9):1524–1538.
- Hull, C. W. 1986. "Apparatus for Production of Three-Dimensional Objects by Stereolithography". Technical Report 4575330, UVP Inc, Upland, CA, Patent and Trademark Office, Washington, DC, U.S.
- Jin, Y., S. Qin, and Q. Huang. 2015. "Out-of-Plane Geometric Error Prediction for Additive Manufacturing". In *2015 IEEE International Conference on Automation Science and Engineering (CASE)*. August 24th-28th, Gothenburg, Sweden, 918–923.
- Law, A. C., C. C. R. Wang, J. Kucukdeger, E. and Z. Kong. 2022. "Process Parameter Optimization for Reproducible Fabrication of 3D-Printed Tissue Scaffold Porosity". *Biofabrication*. publication pending.
- Lin, T.-Y., P. Goyal, R. Girshick, K. He, and P. Dollár. 2017. "Focal Loss for Dense Object Detection". Technical Report abs/1708.02002, Computing Research Repository, Cornell University, Ithaca, New York, USA.
- Mikhail, E. M., J. S. Bethel, and J. C. McGlone. 2001. *Introduction to Modern Photogrammetry*. 1st ed. Hoboken, New Jersey, U.S.: Wiley.
- Pfuhl, P., and M. Degünther. 2021. "Compensation of the Position Errors of Optical Elements by Adapting Their Additively Manufactured Mounting Structure". In *Proceedings Volume 11876, Optical Instrument Science, Technology, and Applications II*. September 12th, virtual, online, 1–7.
- Shen, Z., X. Shang, Y. Li, Y. Bao, X. Zhang, X. Dong, L. Wan, G. Xiong, and F.-Y. Wang. 2019. "PredNet and CompNet: Prediction and High-Precision Compensation of in-Plane Shape Deformation for Additive Manufacturing". In *2019 IEEE 15th International Conference on Automation Science and Engineering (CASE)*. August 22th-26th, Vancouver, BC, Canada, 462-467.
- Shen, Z., X. Shang, M. Zhao, X. Dong, G. Xiong, and F.-Y. Wang. 2019. "A Learning-Based Framework for Error Compensation in 3D Printing". *IEEE Transactions on Cybernetics* 49(11):4042–4050.
- Wang, A., S. Song, Q. Huang, and F. Tsung. 2017. "In-Plane Shape-Deviation Modeling and Compensation for Fused Deposition Modeling Processes". *IEEE Transactions on Automation Science and Engineering* 14(2):968–976.
- Wang, R., A. C. Law, D. Garcia, S. Yang, and Z. Kong. 2021. "Development of Structured Light 3D-Scanner with High Spatial Resolution and its Applications for Additive Manufacturing Quality Assurance". *The International Journal of Advanced Manufacturing Technology* 117(3):845–862.
- Wi, K., V. Suresh, K. Wang, B. Li, and H. Qin. 2020. "Quantifying Quality of 3D Printed Clay Objects Using a 3D Structured Light Scanning System". *Additive Manufacturing* 32(1):1–13.
- Zhang, S., Z. Sun, J. Long, L. Chuan, and B. Yun. 2018. "Dynamic Condition Monitoring for 3D Printers by Using Error Fusion of Multiple Sparse Auto-Encoders". *Computers in Industry* 105(1):164–176.
- Zhao, M., G. Xiong, X. Shang, C. Liu, Z. Shen, and H. Wu. 2019. "Nonlinear Deformation Prediction and Compensation for 3D Printing Based on CAE Neural Networks". In *2019 IEEE 15th International Conference on Automation Science and Engineering (CASE)*. August 22th-26th, Vancouver, BC, Canada, 667-672.

AUTHOR BIOGRAPHIES

BENJAMIN STANDFIELD is a Master of Science student in the Department of Computer Science and a PhD student in the Grado Department Industrial and Systems Engineering at Virginia Tech. His research interests include process improvements for additive processes utilizing machine learning and analytics as well as intelligent manufacturing systems architectures. His e-mail address is bsta7599@vt.edu.

RONGXUAN WANG is a PhD student at the Grado Department of Industrial and Systems Engineering at the Virginia Tech in Blacksburg, Va., specializing in sensing techniques such as three-dimensional scanning, digital image correlation, radiology, and thermography. His research focuses on process monitoring and quality control in smart additive manufacturing. His email address is rxwang@vt.edu.

DENIS GRAČANIN received the BS and MS degrees in Electrical Engineering from the University of Zagreb, Croatia, in 1985 and 1988, respectively, and the MS and PhD degrees in Computer Science from the University of Louisiana at Lafayette in 1992 and 1994, respectively. He is an Associate Professor in the Department of Computer Science at Virginia Tech. His research interests include extended reality, Internet of Things and distributed simulation. He is a senior member of ACM and IEEE and a member of AAAI, APS, ASEE, INSTICC, and SIAM. His email address is gracanin@vt.edu.

ZHENYU KONG is received the BS and MS degrees in Mechanical Engineering from Harbin Institute of Technology, in 1993 and 1995, respectively, and the PhD degree in Industrial and Systems Engineering from the University of Wisconsin-Madison in 2004, respectively. He is a Professor in the Grado Department of Industrial and Systems Engineering and his research interests include real time sensing and smart manufacturing. His email address is zkong@vt.edu.