

AUTOSURF: AUTOMATED EXPERT-GUIDED MESHING WITH GRAPH NEURAL NETWORKS AND CONFORMAL PREDICTIONS

Amin Heyrani Nobari^{1,*}, Justin Rey², Suhas Kodali², Matthew Jones², Faez Ahmed¹

¹Massachusetts Institute of Technology, Cambridge, MA

²Lincoln Laboratory, Massachusetts Institute of Technology, Lexington, MA

ABSTRACT

Computational Fluid Dynamics (CFD) is widely used in different engineering fields, but accurate simulations are dependent upon proper meshing of the simulation domain. While highly refined meshes may ensure precision, they come with high computational costs. Similarly, adaptive remeshing techniques require multiple simulations and come at a great computational cost. This means that the meshing process is reliant upon expert knowledge and years of experience. Automating mesh generation can save significant time and effort and lead to a faster and more efficient design process. This paper presents a machine learning-based scheme that utilizes Graph Neural Networks (GNN) and expert guidance to automatically generate CFD meshes for aircraft models. In this work, we introduce a new 3D segmentation algorithm that outperforms two state-of-the-art models, PointNet++ and PointMLP, for surface classification. We also present a novel approach to project predictions from 3D mesh segmentation models to CAD surfaces using the conformal predictions method, which provides marginal statistical guarantees and robust uncertainty quantification and handling. We demonstrate that the addition of conformal predictions effectively enables the model to avoid under-refinement, hence failure, in CFD meshing even for weak and less accurate models. Finally, we demonstrate the efficacy of our approach through a real-world case study that demonstrates that our automatically generated mesh is comparable in quality to expert-generated meshes, and enables the solver to converge and produce accurate results. The code and data for this project is made publicly available at <https://github.com/ahnobari/AutoSurf>.

Keywords: CFD, Mesh, Deep Learning, Graph Neural Networks, Conformal Predictions, GNN, Mesh Segmentation, 3D model segmentation

*Corresponding author: ahnobari@mit.edu

1. INTRODUCTION

Computational Fluid Dynamics (CFD) has revolutionized how engineers analyze fluid flow phenomena, leading to remarkable advances in engineering applications, including aerospace, automotive, and environmental engineering. Accurate and reliable simulations of fluid flow phenomena can lead to safer and more efficient designs, reducing the costs of development and improving overall performance. This can positively impact the broader society by enhancing energy consumption efficiency, reducing environmental pollution, and improving transportation safety. Fast CFD simulations may also lead to the advancement of scientific knowledge in fluid dynamics, leading to further breakthroughs in the field. Faster CFD also enables researchers to perform high-fidelity CFD simulations rather than relying on low-fidelity and less accurate simulations, which is often the case in many computational design approaches for applications such as aerodynamic design [1–3], turbine design [4], and much more. However, accurate simulations of fluid flow require a proper representation of the geometry and the physics of the problem.

The geometry is usually represented as a discrete set of elements, called a mesh, which defines the computational domain and discretizes the equations governing the fluid flow. Therefore, the quality of the mesh has a significant impact on the accuracy and reliability of the CFD simulations. Meshing is the process of creating a finite element mesh. The mesh quality directly affects the accuracy and efficiency of the numerical simulations. While highly refined meshes may seem like a logical approach to ensure precision, this method comes with enormous computational costs. Ideally, the mesh should have enough resolution to capture the details of the fluid flow but, at the same time, should be sufficiently coarse to minimize the computational cost. Therefore, optimal mesh refinements need to be made, typically in a few critical regions of a structure (e.g., an aircraft's nosecone) that impact the fluid flow around it. However, identifying where to apply refinements to generate a good mesh is an intricate task

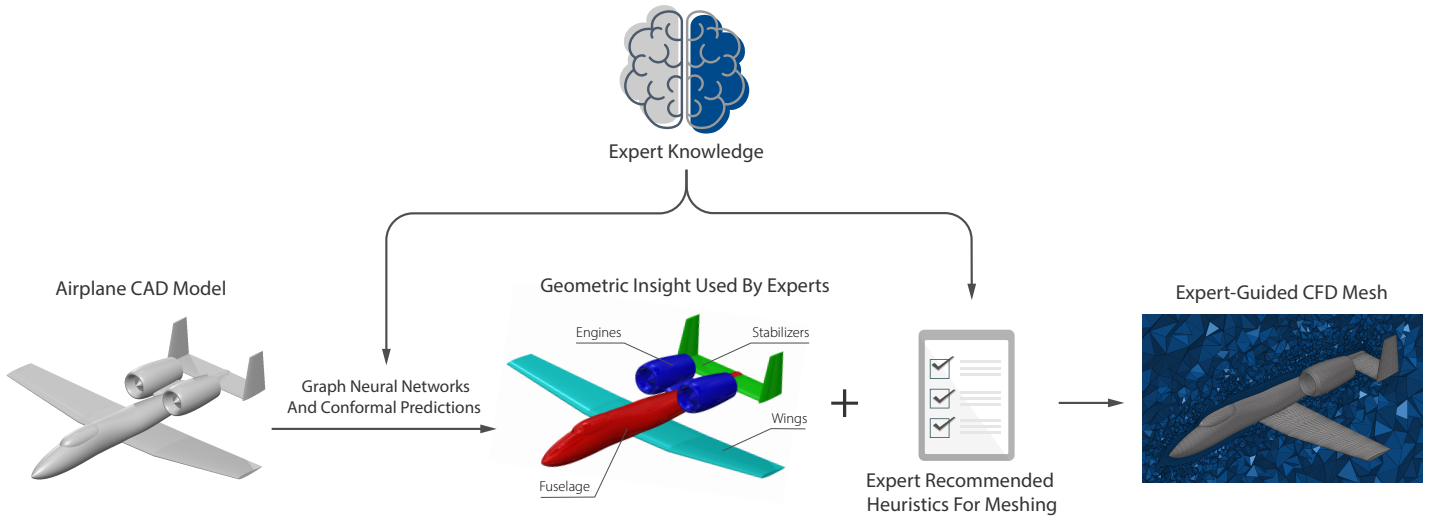


FIGURE 1: EXPERT-GUIDED MESH GENERATION USING GRAPH NEURAL NETWORKS AND CONFORMAL PREDICTIONS

that depends on the specific geometry and flow conditions. This means that significant time and effort are given to the meshing process as it is crucial to the success of a simulation. In case of a poor mesh quality, the entire time-consuming simulation process needs to be repeated. Consequently, generating good computational meshes often requires years of experience from CFD experts. Thus automating the meshing process is of significant value to the overall CFD process and can save significant time and effort and lead to a faster and more efficient design process.

Meshing has been an active area of research for several decades, and a vast number of mesh generation techniques have been proposed. An automated meshing algorithm must generate a mesh that satisfies both accuracy and computational efficiency requirements. In automating mesh generation for CFD, significant research has been conducted on adaptive re-meshing methods [5, 6], where an initial mesh is iteratively refined concurrently with the solver by tracking errors in the domain and refining the mesh based on heuristic rules or optimization [5–11]. However, these iterative adaptive meshing approaches are often slow and incur a high computational cost, making these methods often lead to much longer simulation times which makes them inferior to expert generated meshes that may not be as optimal as adaptive ones (maybe slightly over refined) but will lead to a valid solution in less time.

Despite the progress made in this field, meshing remains a challenging task, especially when dealing with complex geometries. In recent years deep learning models have shown great promise in many different domains of engineering design [12]. Substantial works have been proposed for improving iterative methods such as adaptive re-meshing using deep learning models [13–15]. However, most of these approaches focus on accelerating adaptive re-meshing and ignore the human experts involved in the mesh generation process. Our discussions with CFD experts at the Lincoln Laboratories revealed that, in reality, experts often use specific heuristics or thumb rules to generate meshes instead of adaptive meshing, which can be time-consuming. The rules are specifically handy when the underlying geometries do

not change significantly from one instance to another. Many of these heuristic rules require a deeper understanding of the geometry of the problem. In this paper, we attempt to accelerate this expert-guided meshing process.

We present a machine learning-based scheme for automating the expert-guided mesh generation process for airplane models, with the aim of achieving expert-level mesh quality (Fig. 1). Our approach utilizes Graph Neural Networks (GNN) to gain a deeper understanding of geometry and develop heuristic-based algorithms that automatically generate high-quality meshes. We begin with rough surface meshes of airplane models generated by CAD software and use a GNN-based model for mesh segmentation of airplane parts. In the aerospace industry, reliable risk assessment and uncertainty quantification are essential components of design. To address this, we introduce a component to our approach that utilizes conformal prediction methods [16–18] to quantify uncertainties associated with our model and provide a mechanism for risk handling with marginal statistical guarantees. By enabling designers to set an acceptable risk level, our approach ensures that the generated mesh is, at worst over-refined, with under-refinement occurring only at an acceptable level of risk. The last step of our approach involves projecting the model predictions onto the CAD surfaces to identify the aircraft’s individual parts. Expert-guided heuristic rules are then applied to determine the necessary mesh fidelity and refinements for each part, which are used to automatically generate CFD meshes for each aircraft model (figure 1). We also release a new annotated dataset of aircraft models, which is created by a new augmenting approach proposed in this paper.

Our research paper provides several significant contributions to the field, including:

1. **Application Impact:** We establish an approach for an expert-guided CFD meshing process for aircraft by combining graph-neural networks based predictions with a rule-based approach.
2. **Method:** We introduce a new 3D segmentation algorithm

that outperforms two state-of-the-art models, PointNet++ and PointMLP, for surface classification.

3. **Robust Risk Management:** We present a novel approach to project predictions from 3D mesh segmentation models to CAD surfaces using the conformal predictions method, which provides marginal statistical guarantees and robust uncertainty quantification and handling.
4. **Dataset:** We provide a dataset of realistic airplane models with segmentation labels, which can be used to benchmark future 3D segmentation models.
5. **CFD Case Study:** Through a real-world case study, we demonstrate that our automatically generated mesh is comparable in quality to expert-generated meshes, and enables the solver to converge and produce accurate results.

2. BACKGROUND AND RELATED WORKS

Our approach is based on 3D model segmentation where we take meshes of aircraft models and predict which part of the aircraft each face in the mesh belongs to. We then map these predictions to CAD surfaces, which are then used for determining mesh settings. In this section, we introduce a brief background on 3D segmentation models and conformal predictions, both of which play a major role in our work.

2.1 3D Segmentation and Mesh Segmentation

The popularity of deep learning-based methods has led to the proposal of various learning techniques in the computer vision and computer graphics communities for 3D shape classification and segmentation. Although conventional algorithms exist for this purpose, we will only focus on deep learning-based methods, as they significantly outperform conventional ones.

There are three primary approaches for 3D shape segmentation: 1) Voxelizing 3D shapes and using 3D convolutional neural networks (CNN) to generate 3D semantic segmentation labels, 2) Using point cloud representations and applying point-based approaches to classify each point in the point cloud, 3) Using a surface mesh representation of 3D shapes and classifying either vertices, edges, or faces of the mesh. Our method uses a mesh representation and classifies the faces of the mesh. Mesh is the most physically meaningful segmentation, with faces representing surfaces of the 3D model that would belong to different parts of an object.

In the following sections, we will briefly discuss the state of the art in each of the three approaches to shape classification and highlight some relevant parts of these models that inspired our method.

2.1.1 CNN Based Methods. One straightforward approach for using CNNs in mesh segmentation is to leverage rendered images of 3D objects from different viewpoints and apply existing CNNs that are highly effective in image processing tasks to make predictions on such images. Su et al.[19] were pioneers in this approach, using a multi-view CNN for shape classification, but their method could not be easily adapted for semantic segmentation. Later research developed a more comprehensive

multi-view framework [20] for shape segmentation, where segmentation maps were predicted for each view using CNNs, and the resulting image-level predictions were projected onto 3D models by enforcing label consistency using Conditional Random Fields (CRF) [20].

By converting a three-dimensional shape into a binary voxel format, it is possible to create a grid representation similar to that of a two-dimensional image. This approach enables the same CNN operations used for images to be easily extended to three-dimensional grids, allowing for a seamless transfer of traditional image-based techniques to the realm of shapes. Wu et al. [21] were the first to explore this idea for 3D shapes, and subsequent works have expanded on this approach [22–26]. However, volumetric representations and image-based methods demand significant computational resources and extensive memory usage. Furthermore, binary voxel representations do not readily map to more accurate representations of 3D shapes, such as meshes. Given these limitations, our work departs from image-based and volumetric representations as well as CNN-based methods.

2.1.2 Point cloud Approaches. Point clouds are a common representation method of 3D shapes that can be easily obtained from various other types of 3D shape representations, leading to significant research efforts in utilizing point-based deep learning methods for 3D shape analysis. One of the seminal works in this area is PointNet [27], which utilizes 1×1 convolutions on point cloud coordinates for 3D shape prediction, as well as introducing the Transformation Network (T-Net) to enable model invariance to geometric transformations such as rotations and translations [27]. This concept has proven to enhance learning generalization and can be applied to most representations using point/vertex positions, such as point clouds or meshes. Our approach also employs a variant of T-Net to take advantage of these benefits.

PointNet++ [28], a follow-up to PointNet, partitions points into groups to better capture local structures, resulting in significant performance improvements [28]. Other approaches incorporate information from local neighborhoods to perform dynamic updates by computing the similarity between points based on their euclidean distance in the feature space [29]. Some models, such as RSNet [30], sort points and treat them as sequences, using Recurrent Neural Networks (RNN) to capture local features. PointCNN [31] expands the convolution concept beyond local grids to a χ -convolution that operates on points located within their respective Euclidean neighborhoods. However, PointCNN is not invariant to point permutation, which is addressed in PointConv [32], extending the notion of convolution to an operator invariant to order.

One of the current best approach, PointMLP [33], proposed at the International Conference on Learning Representations (2022), removes convolution and instead employs simple multi-layer perceptrons (MLP) as operators. It extends the ideas of PointNet and PointNet++ and outperforms all other models in segmentation and classification tasks, including those using mesh representations [33]. In the results section, we show how our proposed approach outperforms PointMLP.

Most of these approaches capture local features by hierarchically grouping points and learning group-wise features. However,

if the mesh structure is available, its edges (i.e., adjacency information) can enable local feature extraction without the need for grouping and complex operations. This is the topic of the next section and the primary reason for our use of mesh representations.

2.1.3 Mesh Based Approaches. The non-uniform polygonal mesh serves as the foundation for 3D data representation in computer graphics. This approach utilizes a smaller number of large polygons to cover expansive, flat regions, while employing a larger number of polygons to capture intricate details. A mesh accurately captures complex structures and distinguishes them from nearby surfaces by explicitly representing the surface topology. As such, meshes are considered one of the most accurate representations of 3D data. Models that utilize meshes directly are able to capture local features as needed, eliminating the issue of locality present in point clouds. Additionally, mesh representations do not suffer from permutation issues since they capture adjacency rather than a set of unordered points, as seen in point clouds. These advantages make meshes the most efficient and promising representation of 3D data for learning.

Numerous approaches have been explored for mesh-based deep learning models, including the popular MeshCNN approach introduced by Hanocka et al. [34]. This model utilizes mesh edges as the basis for its approach and computes features for each edge in the mesh, which are invariant to transformations. Convolution is then applied to these features in the neighborhood of each edge, determined by the edges sharing faces in the mesh. Despite its success, MeshCNN has limitations due to its reliance on the convolution operator, making it necessary to refine or coarsen meshes to a specific number of edges, thereby limiting its broad applicability.

To overcome this limitation, a more generalizable and versatile approach has been recently explored, namely the graph-based approach. In this class of models, graph neural networks (GNNs) are utilized to analyze 3D meshes, which lend themselves to the graph representation, with vertices serving as nodes. As such, GNNs are well-suited to handle meshes, and recent works have focused on this approach [35–38]. We combine the GNN-based approaches with point cloud-based approaches in our work to present a new architecture that encompasses the benefits of both classes of models. Further details of our methodology are discussed in the following sections. Next, we shift our focus to a new method, which allows us to make predictions with provable guarantees on the error rates.

2.2 Conformal Predictions

Conformal prediction is a framework in machine learning that provides a measure of confidence in the predictions made by a model without any assumptions on the underlying predictive model. It is based on the idea of constructing prediction regions around each point in the input space rather than just making a single-point prediction. These prediction regions can be calibrated to provide a desired level of confidence. This is made possible through the principle of exchangeability of the data (i.e., the order of the data does not affect the predictions).

Let’s consider a classification task where we have n data-points $\{(X_i, Y_i)\}_{i=1}^n$ with features $X_i \in \mathbb{R}^f$ and a classification

label $Y_i \in \mathcal{Y} = \{1, 2, \dots, C\}$ with C possible discrete classes in the label space. If we train a typical machine learning classifier that outputs a softmax probability distribution for each class, we would have a vague notion of the model’s certainty as these values are just guesses that the model has made. The reality is that there is no guaranteed accuracy/uncertainty that these distributions are reflective of the true distribution P_{XY} . If the cost of mistakes isn’t high, one can report the single most likely class for any input which comes with no statistical guarantees. However, in our work, each CFD simulation can be time-consuming, and knowing a measure of confidence in the machine learning algorithm can reduce mistakes. When we obtain predictions on mesh faces, we project them onto CAD surfaces. Each CAD surface receives information from many mesh faces associated with it, and given the mesh CFD fidelity is determined by these predictions, it is important to have some statistical guarantees on the accuracy of the predictions. Conformal predictions enable this by constructing prediction sets $\hat{\mathcal{C}}_\alpha \subseteq \mathcal{Y}$ with some risk level α for each input rather than a single prediction. Most notably, these prediction sets come with a marginal guarantee of obeying:

$$\mathbb{P}[Y \in \hat{\mathcal{C}}_\alpha(X)] \geq 1 - \alpha \quad (1)$$

These approaches enable such a guarantee by running the model on a test set of data that the model has not seen before and calibrating the uncertainty outputs of the model to a threshold that the model must meet for the above statistical guarantee [17]. This kind of statistical guarantee gives us a much better insight into the uncertainty of the predicted label distributions of the deep learning models which are arbitrary guesses with no . For these kinds of classification problems, many researchers have proposed conformal prediction schemes which can be easily applied to machine learning models [16–18]. In our works we take the approach of Adaptive Prediction Sets (APS) proposed by Romano et al. [17] to construct predictions sets which we then use for a voting algorithm.

Making incorrect predictions for CFD meshing can significantly delay the meshing process. By providing a measure of confidence, conformal prediction can help CFD experts make informed decisions and take appropriate action to prevent failures or minimize risks. This is particularly crucial in simulations that may require very large computational resources, which may not be readily accessible at all times. In such cases where failure cannot be tolerated, conformal predictions can be applied with very low-risk acceptance, while in cases where failure can be tolerated to some extent, this risk can be relaxed in favor of a less conservative mesh which can speed up the simulation.

3. METHODOLOGY

In this paper, we introduce a new model for mesh segmentation which is designed to work with very small datasets without overfitting while having a complex and deep architecture that takes aspects of GNN-based models as well as point cloud-based models to construct a hybrid model with greater accuracy. The overall architecture of our model can be seen in figure 2. We then propose an approach to take the predictions on the surface mesh and project them to the CAD surfaces, which can then be

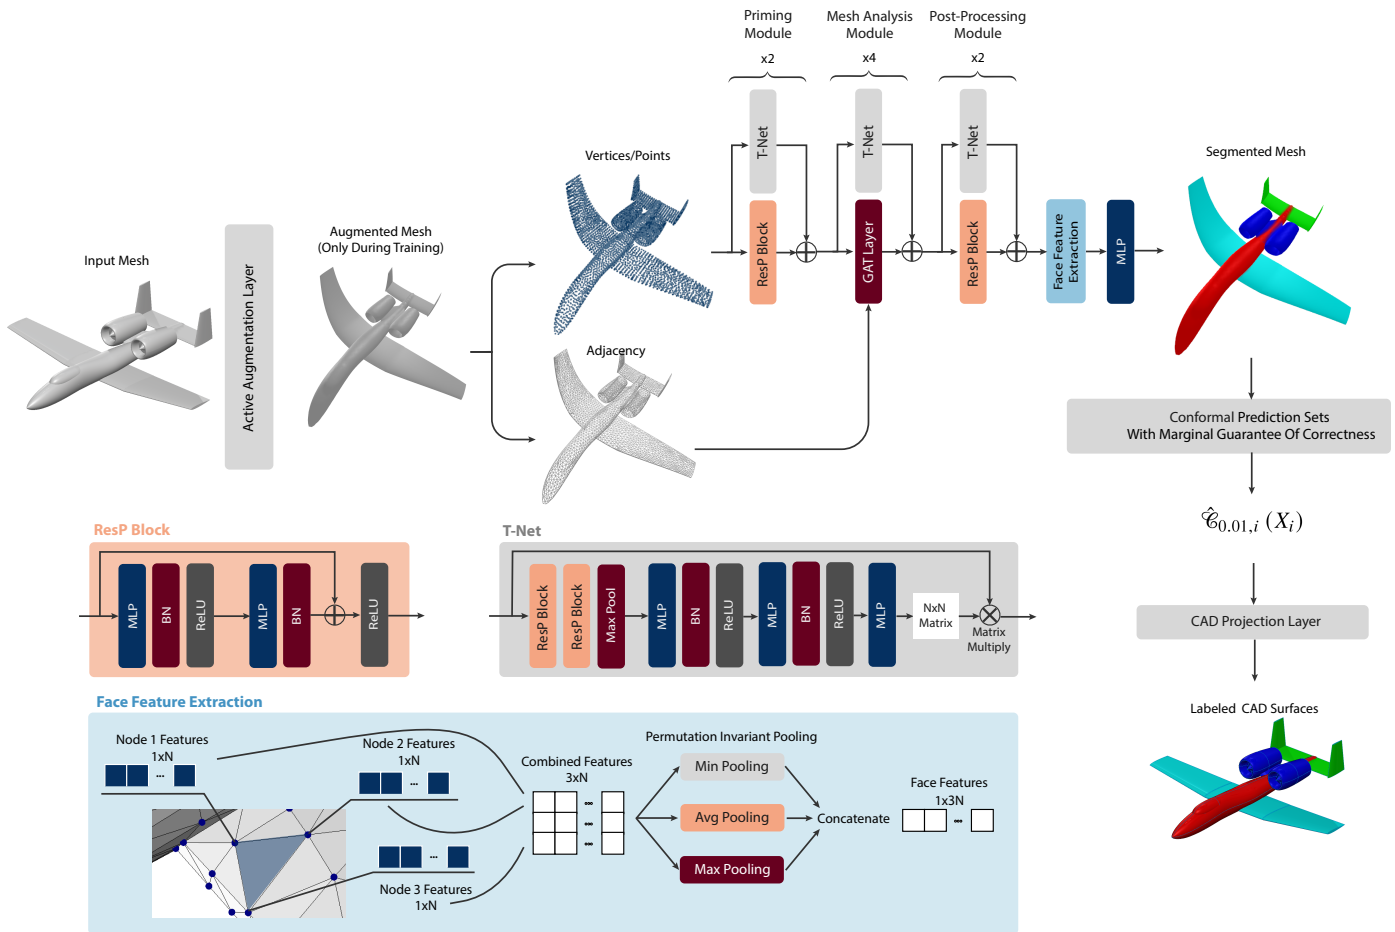


FIGURE 2: THE PROPOSED MESH SEGMENTATION MODEL IS A COMBINATION OF POINT-BASED AND GRAPH NEURAL NETWORK (GNN)-BASED APPROACHES. DURING TRAINING, EACH INPUT FIRST GOES THROUGH AN AUGMENTATION LAYER. THEN THE AUGMENTED MESH IS PROCESSED THROUGH TWO RESP BLOCKS, FOLLOWED BY FOUR GAT LAYERS AND ANOTHER TWO RESP BLOCKS, WITH A T-NET RESIDUAL APPLIED AT EACH LAYER. THE RESULTING FEATURES FOR EACH FACE ARE EXTRACTED AND PASSED THROUGH A CLASSIFIER MLP TO GENERATE SEGMENTATION LABELS FOR THE MESH. ONCE THE MESH IS SEGMENTED, MESH FACES ARE PROJECTED ONTO CAD SURFACES BASED ON CONFORMAL PREDICTION SETS, WHICH PROVIDE MARGINAL STATISTICAL GUARANTEES OF ACCURACY.

used to construct the final CFD mesh. In the sections that follow we will describe the details of our approach.

3.1 Mesh Segmentation Approach

In our model, we focus on developing large models that can be trained on very small datasets. This is particularly important in many engineering design applications as usually the size of datasets in engineering design is very small, which makes the application of large deep learning models difficult [12]. However, it has been demonstrated that proper data augmentation can overcome this challenge and enable the application of large deep-learning models without overfitting [12, 39, 40]. Given this, we introduce an active augmentation layer to our model to enable better generalization of learning and training with few samples. We describe in detail what this layer does and how it is implemented in training in later sections.

After the augmentation, we introduce a novel architecture (figure 2) for mesh face classification using a combination

of point-based and graph-based deep learning approaches which combine the effectiveness of point-based models and the versatility of mesh graph-based models to allow for easy and accurate mesh segmentation.

Finally, we use the conformal predictions approach known as Adaptive Prediction Sets (APS) proposed by Romano et al. [17] to construct prediction sets from our segmentation model's output such that a marginal statistical guarantee can be given to the users when it comes to the possible labels each face in the mesh may be associated with. In the following sections, we describe each of the three aspects of our approach in detail.

3.1.1 Scheduled Augmentation Layer. When working with small datasets, large deep learning models tend to overfit the training data and fail to generalize to data the model has not been trained on. This leads to a lack of generalization which makes the models unreliable in real-world applications. There are several ways to overcome this problem, one way is to simply gather more data, but that may be too expensive or even imprac-

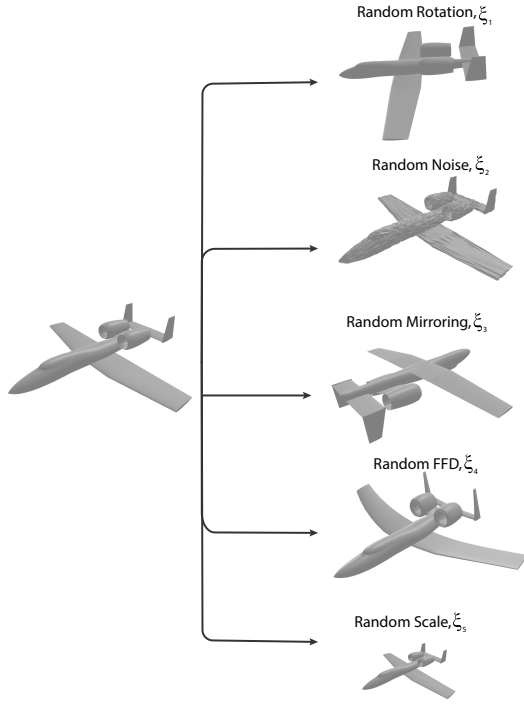


FIGURE 3: VISUAL DEMONSTRATION OF THE FIVE ACTIVE AUGMENTATION APPROACHES WE EMPLOY IN OUR AUGMENTATION LAYER.

tical. Another way to overcome this issue is to take the input data and alter the data slightly so as to prevent the model to overfit or memorize the data. In computer vision and image data, this is typically done by applying random transformation and noise to the data [39]. Here we propose an augmentation layer that applies random changes to the input data to prevent overfitting. In our approach, we introduce a similar method but for 3D data. Our contribution to the augmentation layer is, therefore, two-fold. First, we introduce several augmentation techniques that we find to improve the generalizability of our model. Our second contribution is how we change the extent of our augmentation while the training is going on. In this section, we will first discuss the details of our augmentation pipeline/layer, then we will describe the dynamic augmentation mechanism we use during training.

Augmentation of 3D mesh data: Drawing inspiration from image augmentation methods [39], this paper proposes augmentation mechanisms for 3D mesh data to prevent overfitting in deep learning models. The proposed mechanisms include random rotations in all three directions, adding noise to data by applying random noise to the position of vertices in the mesh, mirroring 3D models with respect to different planes, random application of Free-Form Deformations, and random scaling of the model. One of the common approaches for image augmentation is the random rotation of the images, which is easy to adapt to 3D data by randomly rotating 3D objects in space. Unlike image-based approaches that apply one rotation angle, we apply random rotations in all three directions. Another common augmentation technique is adding noise to data, which we also use in our layer

by applying random noise to the position of vertices in the mesh. Another augmentation mechanism we employ is mirroring 3D models with respect to the XY , XZ , or YZ planes randomly. It also observed that warping images could be an effective form of augmentation in computer vision. We translate that idea to 3D, by using a random application of Free-Form Deformations (FFD) as analogous to the wrapping technique used for images. Finally, we have random scaling of the model, which either enlarges or shrinks a model slightly. These augmentation mechanisms are all demonstrated in Fig.3 to help visualize how they work. However, as we discussed earlier, we intend to control these augmentation mechanisms during training. As such, each augmentation mechanism will be controlled by a specific parameter.

For rotation, we randomly sample a 3D angle, Θ , from a uniform distribution of $\Theta \sim U(0, \xi_1) \in \mathbb{R}^3$ and the control parameter ξ_1 determines the extent of augmentation. Similarly, for the noise, we sample random noise for each vertex (v_i) from a uniform distribution $\epsilon_i \sim U(0, \xi_2) \in \mathbb{R}^3$ and the control parameter ξ_2 determines the extent of augmentation. For mirroring, we apply the mirroring across each of the 3 XY , XZ , and YZ planes randomly with a probability of $0 \leq \xi_3 < 1$, which controls how likely the mirroring is to occur during augmentation. For warping, we apply the FFD by deforming the model randomly in all directions, and we randomly sample the extent to which these deformations occur from a uniform distribution $U(-\xi_4, \xi_4)$, where ξ_4 determines how aggressively the model is deformed. Finally, for scaling, we apply a uniform scale adjustment across all three dimensions, which is sampled from a uniform distribution $SA \sim U(-\xi_5, \xi_5)$, and the model is scaled by $1 - SA$. The ξ_1 , ξ_2 , ξ_3 , ξ_4 , and ξ_5 parameters control how aggressively the models are augmented which allows us to control the extent of augmentation during training. More details can be found in the authors' code and data, which will be made available upon acceptance of the paper.

Augmentation Schedule: When training starts, we start with no augmentation and based on the target value of the ξ_1 , ξ_2 , ξ_3 , ξ_4 , and ξ_5 parameters the augmentation increases as training goes on. In general, we have an augmentation schedule that increases the extent of augmentation as every epoch goes on. The following equation describes this schedule:

$$\xi_i^T = \xi_i^\infty \times \frac{(1 + \cos(\frac{T\pi}{\tau} - \pi))}{2} \quad (2)$$

where T is the current epoch of training and ξ_i^∞ is the target value of the augmentation parameter while ξ_i^T is the adjusted augmentation parameter at epoch T and τ is the scaling factor which is set to the maximum number of epochs (the total number of epochs that we plan to train the model).

While training the deep learning model, we increase the level of augmentation gradually until the model no longer overfits the data. To ensure this, we keep track of the validation accuracy at the end of each epoch and save checkpoints of the model's weights when the accuracy is maximized. However, we do not want to make the task too difficult for the model, so we simultaneously decrease the learning rate of the optimizer as the level of augmentation increases. This helps the model adjust smoothly

to the more challenging task and prevents it from straying too far from the optimal parameters it has learned. Thus, we strike a balance between augmentation for generalizability and not hindering the model’s learning. More implementation details can be found in our code and data, which will be available to the public upon acceptance of this paper.

3.1.2 Mesh Segmentation Model. As discussed in previous sections, point cloud-based models have shown great success in 3D shape segmentation and classification, even outperforming the latest and best mesh-based and GNN-based models [33]. However, these models neglect the most accurate and flexible representation of 3D data [34], which is the mesh. Instead, they use neighborhood grouping and sampling to extract local relationships. Despite this, even when specialized for specific tasks, the latest point cloud-based model PointMLP [33] is comparable in performance to GNN-based models [35, 36]. Therefore, we cannot overlook the performance advantages of point cloud-based methods despite their limitations. To create a hybrid model that combines the strengths of both, we incorporate the lessons learned from point cloud-based models into GNN approaches, aiming to achieve the benefits of point cloud models and the versatility of GNN models.

Figure 2, shows the overall architecture of our model. There are several stages in our architecture, which we discuss in more detail here.

Priming Module: As it can be seen, the model starts with a priming module which includes an adaptation of the pointMLP’s Residual Point (ResP) [33] block with the addition of T-Nets proposed by PointNet [27]. As we mentioned in the background section, the T-Net is a mechanism for transformation invariance, which is important when it comes to 3D data, and this matter was taken care of in PointNet++ [38] and PointMLP [33] by grouping and sampling through their geometric affine modules. Since we intend to use the mesh directly in our model and do not wish to perform geometric grouping and sampling, we introduce the T-Net. Our implementation of T-Net is shown in Figure 2. It takes as input a point set and using a few ResP layers, determines a linear transformation matrix for the point set which is then applied to the point set’s features.

In this module, despite only having point-based components, we do not down-sample the vertices in the mesh to a specific number of points, as this would not allow us to apply GNNs to the entire mesh effectively. One of the main limitations of the PointBased models, is that during training, only subsets of all the points will be seen, while in our model, we not only input all the points but make classifications on all of our mesh faces. Another limitation of the point cloud-based model is that predictions will be made at the points, not surfaces. Furthermore, predictions in point cloud-based models are made at the points, while true labels are only meaningful for mesh faces in 3D segmentation. Thus, for point-based models, less accurate labels must be determined for individual points, and combining vertices of the mesh is not feasible, as all the vertices associated with faces may not have been captured during the subsampling of training. Therefore, by inputting all the vertices and using GNNs, we avoid this issue and

properly analyze the entire model.

Mesh Analysis Module: Once the data goes through the point-based priming module, we take the features from this stage as the input features for the mesh analysis module. In the mesh analysis module, we represent our data as graphs with node features from the priming module and adjacency matrix same as the mesh adjacency. In this module, we use a GNN approach called Graph Attention Convolution (GAT) [41]. For the sake of brevity, we will not discuss the details of GAT here and refer the readers to the original body of work for more detail [41]. Simultaneously, we find applying the T-Net from the priming module helps with the model generalization. For the GAT layers, we use 8 attention heads with 64 features each and apply a 10% dropout during training.

Post-processing Module: After the mesh analysis module, the resulting features go through the post-processing module, which is made up of ResP blocks and T-Nets. After that, the features of the mesh vertices are aggregated for each face. That is, the features of the vertices associated with each face are aggregated using an order-invariant face feature extraction function (see figure 2). This aggregation function must be order invariant because if the model sees different permutations of the points associated with a face, the outcome should not change. In our approach, we perform aggregation by min-pooling, average-pooling, and max-pooling and concatenate the three in that order. In this way, the features from all three vertices are captured without loss of information while the order invariance of the aggregation function is maintained. Finally, in the post-processing module, we apply an MLP to each face’s features to predict the label for each face.

During training, we use the categorical cross-entropy loss to train the model for segmentation:

$$L_{CLS} = -\frac{1}{N} \sum_i^N \sum_j^M y_{ij} \log(\hat{y}_{ij}) \quad (3)$$

where L_{CLS} is the classification loss that we minimize during training and N is the number of faces in a batch during training and M is the number of possible labels, and y_{ij} is the true label for the face i and class j while \hat{y}_{ij} is the label predicted by the model. Beyond this loss, we also apply regularization loss to the transformations predicted by every layer of T-Net we have in our model. This is needed as the dimensionality of the features in the model increase, and the predicted transformations become much larger, which leads to overfitting and instabilities in the model. To mitigate this issue, we apply a regularization loss to the transformations predicted by T-Nets:

$$L_{T-reg} = \sum_i^N \|I - A_i A_i^T\|^2 \quad (4)$$

Where L_{T-reg} is the regularization loss minimized and N is the number of T-Net layers and A_i is the transformation predicted by the i -th T-Net layer. The total loss that we minimize during training is, therefore:

$$L = L_{CLS} + \gamma \times L_{T\text{-reg}} \quad (5)$$

Where gamma is the weight given to the regularization loss.

3.1.3 Conformal Predictions. Conformal prediction is a powerful technique for improving the reliability and robustness of machine learning models, making it a useful tool in a wide range of engineering applications. At this point, we have described a standard 3D segmentation model. However, the performance of deep-learning models is dependent on the data, and they simply make point predictions that are best guesses. The predictions for new test cases do not come with any statistical guarantees. This means that if we only base our classification on the model’s best guess, we cannot guarantee any accuracy. In general, uncertainty quantification with any kind of guarantee based on learning is difficult and most work in this area is restricted to Gaussian Processes. However, we could easily overcome this issue in a model-agnostic way by using conformal predictions. In our work, we use the approach proposed by Romano et al. [17] to construct prediction sets that are marginally guaranteed to include the correct prediction 95% of the time. More formally:

$$\mathbb{P} [Y_i \in \hat{\mathcal{C}}_{0.05,i} (X_i)] \geq 0.95 \quad (6)$$

Where $\hat{\mathcal{C}}_{0.05,i}$ is the prediction set for the input sample (mesh face) X_i and Y_i is the true label. We then use these prediction sets to map the predictions on the mesh to the CAD surfaces. One thing to note here is that if the acceptable risk level (i.e., 0.05) were to be lowered, this may lead to the model becoming more conservative, leading to more surfaces having uncertain labels associated with them, which can lead to an over-refined mesh more often, leading to more expensive CFD simulation. As such, great care must be given to defining an appropriate risk level. In this work, we found that for our model a 0.05 risk led to no increase in the number of mistakes by the model but correctly enabled the identification of uncertainties and removed any under-refinements.

3.2 Mapping Mesh Segmentation to CAD Segmentation

In order to generate an effective Computational Fluid Dynamics (CFD) mesh for an airplane model, it is necessary to apply specific mesh settings to different surfaces of the aircraft. As a result, it is not sufficient to classify different parts of the airplane in a single mesh. Instead, each surface of the aircraft must be meshed with specific settings. Thus, the most practical approach to automating the CFD meshing process would involve mapping the classifications from the mesh to the corresponding surfaces in a CAD model. These surfaces can then be meshed in a CFD meshing platform using the appropriate settings for each surface type.

Estimating Mesh to CAD Surface Distance: To map predictions from the mesh to the CAD model, we take the predictions for each face in the mesh associated with the CAD files and calculate the location of the centroid of the face. Then we measure the distance from the centroid of each face to all the CAD surfaces and identify the closest surface in the CAD for each mesh face.

To measure the distance to different CAD surfaces we take the B-Spline surfaces of the CAD file and discretize the surfaces numerically using the control points defining the B-Spline surface and obtain a grid of points in each surface. Then we measure the distance between any given face centroid and all of the points in the discretized surface and take the minimum of all these distances as the distance from a face centroid to any given CAD surface:

$$D_{ij} = \min_k \|x_i - x_{jk}\| \quad (7)$$

Where D_{ij} is the distance between face i and surface j and x_i is the location of the centroid of face i and x_{jk} is the location of the k -th discretized point of the j -th surface.

Voting Mechanism for CAD Surface Mapping: Once the CAD surface for each face is identified, the faces will vote on the nearest CAD surface based on the conformal predictions made previously, which have a 95% marginal correctness guarantee. The votes for each surface are then tallied, and if a single class obtains the majority vote (i.e., 50%+1 vote), it is selected as the surface type. If this does not occur, we select the surface classification from the top two classes with the most votes that require more refined mesh settings (see section 3.3 for details). This approach ensures that the mesh is not under-refined and can resolve the flow, leading to valid solutions. However, it may result in some surfaces being over-refined, which comes at a higher computational cost for the solver. When the conformal prediction sets have a 95% marginal correctness guarantee, the model can accurately predict a singular majority in most cases. If this is not possible, our model takes a more conservative approach to meshing to avoid any negative outcomes. Therefore, with an accurate enough model, we can practically guarantee valid mesh settings.

3.3 Expert-Guided Automated Meshing

After obtaining the CAD surface classifications through voting, we use expert-guided rules to generate meshes automatically for the CFD simulation under specific flow conditions and simulation types (as shown in Figure 1). These rules would be specific to the CFD simulation type and flow conditions. The primary idea in our work is that a database of rules would be established for each kind of simulation type, and equations will be established for specific mesh settings to be adjusted based on flow conditions. In our work, we focus on a specific setup for CFD simulations and use rules established by experts for mesh generation. Specifically, we develop rules for simulations intended to resolve unsteady flow structure in wake regions for aero-optical work. Furthermore, the mesh generated is meant to be simulated in the FUN3D CFD solver (second-order time accurate simulation), with Detached Eddy Simulation (DES) turbulence model. For this work, we limit the scope of work to one specific flow condition as well. This is done to enable us to demonstrate a proof of concept in this paper, and future work will generalize the rules to more situations. The flow conditions we develop our rules around are flight at sea level, at Mach 0.8, and Angle of Attack (AoA) 0°. We also resize airplane models to have a

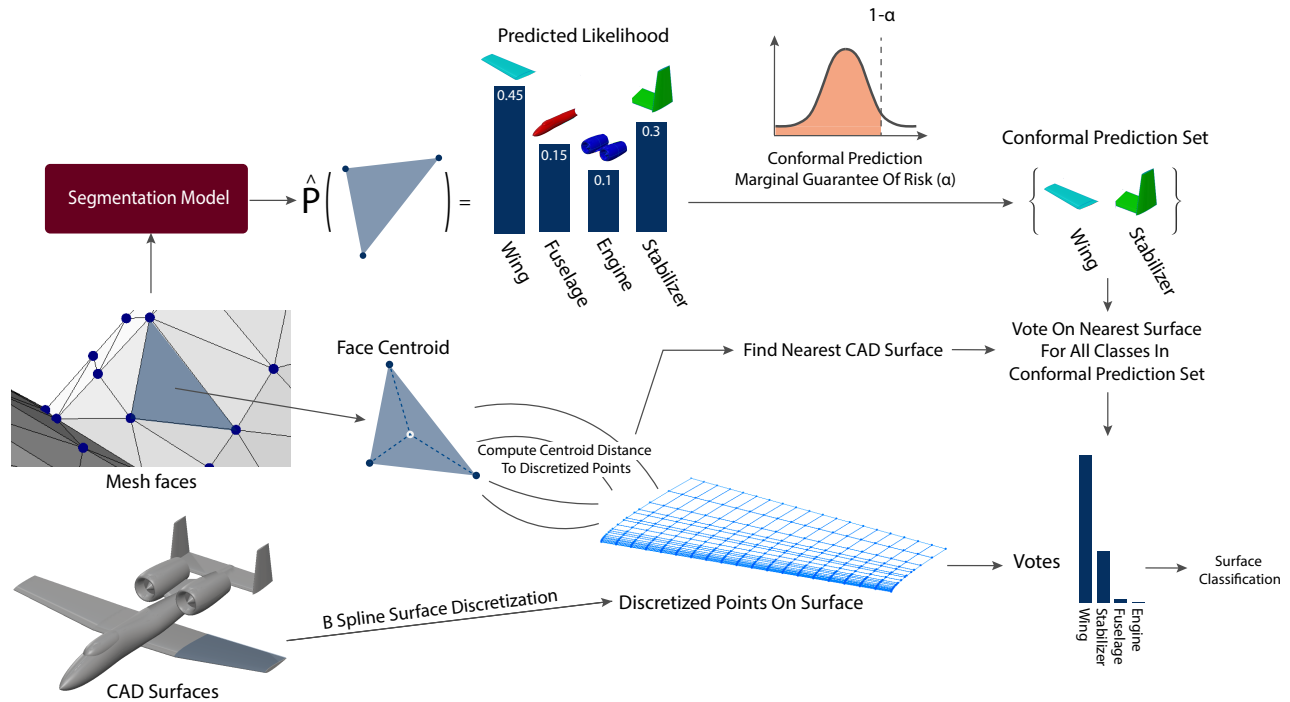


FIGURE 4: OUR APPROACH FOR MAPPING MODEL PREDICTIONS TO CAD SURFACES WITH MARGINAL STATISTICAL GUARANTEES USING CONFORMAL PREDICTIONS.

length of 30m. We then develop rules (based on Pointwise meshing software settings) for different aircraft parts to automate the mesh generation process. We do not go over the specific rules in this section to improve the readability of the paper. These rules are specified in appendix A. Although these rules are established automation of the process is not explicitly discussed further in this paper, as mesh settings databases are established by experts at the Lincoln laboratories, and the sensitive nature of their work limits the amount of information that can be discussed publicly. As such, we only provide the details for the above-mentioned flight conditions and see the process of automation as a black box that involves providing surface predictions to experts who can automate the meshing in their desired meshing software closing the loop of expert-guidance integration (see figure 1).

An important thing to note is that the settings that must be applied to the wing are the most refined, followed by the stabilizers, and finally the fuselage has the least refinements. As such when a majority is not reached in CAD surface classification (see section 3.2) priority is given to the classification that requires the most refined mesh.

4. DATASET

So far, we have discussed a framework for mesh classification and how we can map these predictions to CAD files and use expert guidance to generate CFD meshes. However, there are no publicly available datasets for the purpose of our work. As such to train our model we needed to create a new dataset of airplane models which meet the following criteria. First, the models used for training must be real and accurate aircraft designs and not include examples of toy aircraft and concept designs which are common in datasets such as ShapeNet. This is because, in many

of these cases, the features in the models are not indicative of realistic aircraft designs and therefore, may confuse the model and reduce generalizability to real models, and bias the model towards the features observed in unrealistic samples. The other criterion is that the models in the dataset must have segmentation labels pertaining to different aircraft parts, specifically the wing, stabilizers, fuselage, and engine as these are the parts of the aircraft that are important for mesh settings. The third criterion is that the models must include accurate CAD representations associated with them as well since we need to classify the CAD surfaces to apply the mesh settings properly in meshing software.

To create a model that is practically useful for CFD simulations at Lincoln laboratories, we created a dataset of realistic aircraft using Open Vehicle Sketch Pad (OpenVSP) [42]. OpenVSP, initially developed by NASA, is an open-source tool for creating parametric aircraft geometry that enables the engineering analysis of those models. To create the dataset, we use a few selected aircraft designs from OpenVSP and then apply parametric sampling within OpenVSP and data augmentation techniques outside of OpenVSP. In the following sections, we describe how we went about doing this in a more detailed manner.

4.1 Dataset Backbone

As discussed above, we chose OpenVSP to create a dataset of realistic aircraft models, as their website contains many aerodynamically accurate airplane models within their software, which are made publicly available under the name VSP Hanger¹. However, not all of these models can be used easily, as we will see. Before we can use any given model, we need to ensure that we can

¹<https://hangar.openvsp.org/>




FIGURE 5: THE DATASET IS DERIVED FROM TEN BASE AIRPLANE MODELS, EACH BELONGING TO A DIFFERENT AIRCRAFT TYPE. THE MODELS WERE PROVIDED TO US BY CFD EXPERTS, WHO ENSURED THAT THEY COVER A WIDE RANGE OF AERODYNAMIC DESIGNS.

generate valid geometries for each model and label the different parts of the aircraft. This is a crucial part of the dataset curation, and as a result, we attempt to automate the process as much as possible to enable future dataset expansion.

The initial task is to develop a labeling method for different parts of the aircraft. OpenVSP has a feature that enables users to allocate various parts of the geometry to different sets, which is useful for our goal of segmenting airplane models. We manually assign different parts of the aircraft geometry to various sets, creating lists of sets for each aircraft’s fuselage, wing, stabilizers, and engines. We extract separate meshes for each set in the OpenVSP model and use MeshLab’s [43] boolean operations to combine them, resulting in a water-tight mesh. As we keep track of which OpenVSP set each face in the mesh belongs to during boolean operations, we can generate labeled meshes using this information and our manual labeling. We will not delve into the intricacies of this process since the dataset is not the primary focus of our work. However, we will make the code we use for data processing and the data itself publicly available and refer interested readers to our code and data for further detail.

The final note on the data is that given the process of manual labeling and verification of geometric validity is manual and labor intensive at this moment; our dataset is limited to ten accurate aircraft models selected by CFD experts at Lincoln laboratories, which are displayed in figure 5. Despite being a small number of models, there is significant variation and diversity in the dataset (figure 5) to capture many different types of aircraft designs. However, this small size presents a major challenge when it comes to applying large deep-learning models, as training a model on eight aircraft models (assuming two models are set aside for validation) is almost certainly going to lead to overfitting. To overcome this, we propose an augmentation approach which we discuss in the following section.

4.2 Working With Small Datasets In Data-Driven Models

Since the small size of the dataset does not lend itself to deep learning, we employ some data-level augmentation, which is applied at the data-gathering stage rather than later during training. In our dataset, we employ two strategies for augmentation: design alterations, and mesh refinement alteration.

For design alteration, we use OpenVSP’s python API. First, we manually analyze the design parameters (geometric parameters) in the OpenVSP graphical interface and determine appropriate design parameters that can be altered without breaking the

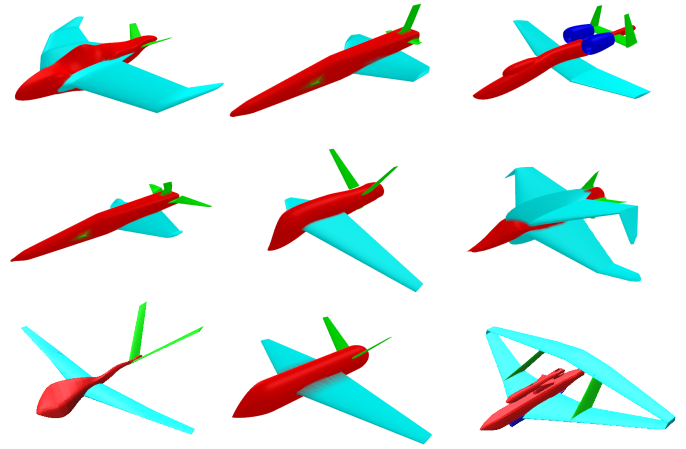


FIGURE 6: SHOWN ABOVE ARE NINE RANDOM AIRPLANE MODELS FROM THE AUGMENTED DATASET, WHICH ILLUSTRATES THE SIGNIFICANT DIVERSITY ACHIEVED BY CHANGING DESIGN PARAMETERS. IT IS IMPORTANT TO NOTE THAT THE EFFECTS OF REMESHING ARE NOT VISIBLE IN THESE IMAGES.

models and the range of acceptable values for said parameters. Once the parameters are determined, we use the OpenVSP API to create randomly altered versions of each of the ten models in our dataset by randomly sampling the values of the design parameters from a uniform distribution of acceptable values. We then generate 20 variations of each of our original models extending the dataset to 210 samples.

One important thing to mention, aside from the limited size of the dataset, is the fact that the mesh generated by OpenVSP has specific patterns in different parts of the aircraft, which could be exploited by a machine learning model to overfit these patterns. Consequently, it is important to ensure the model does not develop a bias towards such patterns. Remeshing the OpenVSP-generated meshes will enable two things for us. The first is that by remeshing at different mesh fidelities, we will augment the dataset and increase the size of the dataset, which is crucial to the success of any learning-based method. The second benefit of this would be that the model cannot be biased toward a specific mesh pattern generated by OpenVSP. To achieve both of these goals, we pick an unbiased uniform remeshing algorithm, isotropic explicit remeshing [43, 44], to remesh each model into five different mesh element sizes and fairly uniform elements, which remove the bias created by software-specific patterns in the mesh. The results of the aforementioned augmentations are visualized in figure 6.

5. RESULTS AND DISCUSSION

So far, we have established the details of our approach. In this section, we report the results of applying the method to the dataset and compare the performance of our model to other existing state-of-the-art approaches. But before we continue, it is important to establish the training details of our model and augmentation layer.

5.1 Model Configuration

To demonstrate the efficacy of our approach in mesh segmentation, we compare our method to two state-of-the-art methods, PointMLP [33] and PointNet++ [38], that outperform most other models or closely match their performance. These models serve as a benchmark for state-of-the-art. In both of these models, to obtain mesh face classifications we simply add the model’s outputs for each vertex on a face and use that as the predictions for the face labels. We also train a naive implementation of a simple graph neural network with six layers of the graph convolution(GCN) [45] layers as the naive baseline. This model serves as a baseline for the simplest possible approach and is trained to demonstrate the notable gap that would exist without better model architectures. For PointNet++ and PointMLP, we use the code provided by the original authors of the paper and only make the final prediction modification without changing anything in their work or model configuration. For details of the naive baseline, we refer you to our code which will be made publicly available upon the acceptance of this paper.

For our model, we use a γ value of 0.1 in our loss (based on Eqn. 5), and for the active augmentation layer, we use target parameters of $\xi_1 = \pi/6$, $\xi_2 = 0.001$, $\xi_3 = 0.2$, $\xi_4 = 0.4$, and $\xi_5 = 0.15$. We find that in our experiments, these values yield the best generalizability for our model configuration. We train our model based on the loss function defined in Eqn. 5 and optimize our model using the Adam optimizer with a learning rate that starts at 10^{-4} and decays 15 times during training at equal intervals and at a rate of 0.65. For further details of the size of layers in our architecture and smaller implementation notes, we refer you to our code which will be made publicly available upon the acceptance of this work. To make the comparisons fair, we train all models for 200 epochs and use the original authors’ code for the above benchmark models.

5.2 Mesh Segmentation Results

As mentioned above, we train all models on our dataset. For training, we use 840 models (based on the first eight in figure 5) and leave 210 models generated based on the last two aircraft (The two models at the bottom right of figure 5) for testing and validation. In figure 7, we visually demonstrate the mesh segmentation predictions made by our model. It is impressive that most of the faces on the meshes are predicted correctly. More importantly, this kind of accuracy will only be amplified when it comes to mapping face predictions to CAD surfaces, as it is much more unlikely that the majority of faces associated with CAD surfaces are incorrectly classified, leading to an even more accurate CAD surface classification.

Besides visual verification of the model’s performance, it is important to quantify the accuracy of the model and compare it to state-of-the-art. As such, we train our model and competing models ten times each and record the means value of the best validation accuracy (percentage of correctly classified faces based on the model’s top guess) for each experiment and model and report the results in Table 1. As evident, our model outperforms all models, including the very recently published and powerful PointMLP, which has established itself as the go-to model for 3D data. Besides that, we do observe that PointNet++ and GCN per-




FIGURE 7: WE PRESENT SIX SAMPLE SEGMENTATION PREDICTIONS GENERATED BY OUR MODEL ON PREVIOUSLY UNSEEN VALIDATION DATA. THE HIGH ACCURACY OF OUR MODEL ON THIS DATASET IS APPARENT FROM THE VISUAL RESULTS. NOTABLY, THE MODEL CORRECTLY CLASSIFIES THE MAJORITY OF FACES ASSOCIATED WITH EACH CAD SURFACE, WHICH ENABLES EVEN MORE ACCURATE CAD SURFACE PREDICTIONS.

form poorly, which is expected as PointMLP is an improvement on the concepts introduced in PointNet++, and the naive GCN model, as expected performs much worse than all other models proving that the baseline of what the simplest possible approach leaves a lot to be desired.

TABLE 1: COMPARISON OF THE MEAN ACCURACY OF DIFFERENT MODELS ON THE VALIDATION DATA. THE VALUES AFTER \pm SHOW THE STANDARD DEVIATION ACROSS TEN TRAINING RUNS.

| Model | Top 1 Accuracy (%) |
|------------------------------------|------------------------------------|
| PointNet++ | 89.14 \pm 0.25 |
| PointMLP | 93.92 \pm 0.15 |
| Naive GCN | 81.23 \pm 1.06 |
| Our Model (w/ Augmentation Layer) | 96.65 \pm 0.12 |
| Our Model (w/o Augmentation Layer) | 92.35 \pm 0.97 |

5.3 Conformal Predictions Supported CAD Surface Classification

For CAD surface classification we take the best-trained model for each method and implement a simple voting algorithm to classify CAD surfaces. This algorithm is very similar to our CAD surface classification (see figure 4) except without conformal predictions such that each face only votes for the top guess of the model. For CAD surface classification we record the number of misclassified surfaces in the validation data for each approach. Furthermore, we calculate the number of misclassifications that would lead to a more refined mesh than necessary (ultimately not failing in simulation) and the number of misclassifications that lead to under-refinement, which can be problematic in simulation. We present our results in table 2. The first observation in these results is the fact that our model which has the most accurate predictions on the mesh as we saw in the prior section, leads to the fewest misclassified surfaces compared to all other models.

TABLE 2: COMPARISON OF THE ACCURACY OF DIFFERENT MODELS ON CAD SURFACE PREDICTIONS ON THE VALIDATION DATA.

| Model | # Incorrect Surfaces | # Under Refined Surfaces | # Over Refined Surfaces | Accuracy (%) |
|-----------------------|----------------------|--------------------------|-------------------------|--------------|
| PointNet++ | 23 | 9 | 14 | 97.26 |
| PointMLP | 13 | 5 | 8 | 98.45 |
| Naive GCN | 42 | 9 | 33 | 95.00 |
| Naive GCN + Conformal | 116 | 0 | 116 | 86.19 |
| Our Model | 7 | 2 | 5 | 99.17 |
| Our Model + Conformal | 7 | 0 | 7 | 99.17 |

The most notable observation is that using the model’s predictions directly without the conformal predictions approach leads to under-refined surfaces, which can affect the simulation efficacy and result in inaccurate solutions, ultimately a failure. This demonstrates the importance of applying conformal predictions to quantify uncertainty and apply the most conservative mesh settings when the model is faced with uncertainty in surface classification. More importantly, conformal predictions enable uncertainty quantification with marginal statistical guarantees, which is a very robust approach, and guarantees under refinement in the mesh will be avoided to a measurable extent which can be set by the user depending on the sensitivity of their work. Another significant observation is that adding the more conservative conformal predictions with 95% marginal guarantee has not caused more over-refinement, which is a testament to the accuracy of our model which is out of the box is good enough that even with the conservative conformal predictions remains accurate. On the contrary, we see that when conformal predictions are applied to models with low accuracies, such as the GCN baseline (see table 2), the conformal predictions fail to provide the necessary guarantee easily, leading to much more over-refinement and an overly conservative CFD mesh. Still, it is impressive that conformal prediction still sees no failed CFD mesh with under-refinement even when the model is inaccurate (See GCN+Conformal in table 2). To the best of our knowledge, this work is the first demonstration of model-agnostic conformal predictions in the engineering design community, and we hope that our results will inspire further research in different engineering applications to improve deep learning model robustness using conformal prediction-based approach.

5.4 Automated Meshing: A Case Study

So far, we have demonstrated that our model is very effective at extracting the necessary geometric insights needed by experts to automate the meshing process. We also discussed the specific settings that can be applied for a specific flight condition. To further showcase the potential of our methods, we conduct a case study on one aircraft model. Specifically, we generate the mesh for the aircraft model resembling the A-10 warthog (top left aircraft in figure 5). We use the expert rules for the flight conditions discussed in section 3.3. The results of the rule-based meshing are shown in figure 8. As we discussed, the flow patterns around the wings are the most intricate and require significant refinement compared to the fuselage, and we can see that given the correct classification, the rule-based meshing successfully refined the mesh around the wing while keeping the mesh




FIGURE 8: THESE IMAGES ILLUSTRATE HOW THE APPLICATION OF DIFFERENT MESHING RULES CAN RESULT IN VARYING LEVELS OF REFINEMENT IN DIFFERENT AREAS OF THE MODEL. SPECIFICALLY, THE TOP TWO IMAGES SHOW HOW OUR MESHING RULES LEAD TO MORE REFINEMENT AROUND THE WINGS, WHILE LESS REFINEMENT IS APPLIED AROUND THE FUSELAGE WHERE FLOW PATTERNS ARE SIMPLER.

refinement lower around the fuselage of the aircraft. In this way, the meshing process for this aircraft can be automated with little input or effort from the users without significant added cost on the solution time. These results show that expert-guided automated meshing for aircraft can be of great value for designers and perhaps provide a viable alternative to very expensive adaptive remeshing approaches or even serve as a great initialization to adaptive remeshing leading to significant time saves in both the meshing process itself and the simulation time.

Finally, to verify that this generated mesh is capable of capturing the necessary details in the simulation, we run the simulation for the generated mesh with the flow conditions mentioned in section 3.3. We observe that the solver converges without any issues, and the resulting simulations capture the details necessary for flow analysis. To showcase the success of the meshing when it comes to the solution, we specifically look at areas of the flow where higher detail must be captured. Specifically, we see that the simulation results have successfully captured the details of the flow in the most complex parts of the flow field near the wings. This shows that expert-guided refinements are crucial for successful flow simulations and accurate results. In the benefit of brevity, we refrain from discussing the CFD details, given that the primary focus of our work is on meshing.

6. LIMITATIONS AND FUTURE WORK

While our methods have been successful, there are areas in which we aim to develop and improve in the future. Currently, our approach is limited to rules based solely on 3D model segmentation, meaning that experts are only able to make rules based on surface types. In the future, we plan to collaborate with CFD experts to create more advanced rules for meshing by identifying other relevant features of 3D models, such as engine inlet and outlet, trailing edges of lift-generating surfaces, and more. This will enable expert guidance to become even more insightful, leading to a more effective and highly optimal mesh generation scheme. We believe that our GNN and point-based model will be able to handle this challenge with ease, and we plan to take this step in the near future.

Another area of improvement for our work would be to expand our dataset to encompass a much larger set of models. This will lead to a more powerful and accurate segmentation model, significantly improving the overall approach. As such, we plan to continue adding models to our dataset in the future and make the data publicly available as a living dataset that will receive updates with more models being added over time.

At this stage, our approach relies on third-party meshing software for CFD mesh generation. However, we believe that a more independent CFD mesh generation pipeline would provide our model with more freedom and enable expert guidance to be significantly more detailed. Therefore, significant effort must be devoted to developing such a pipeline to enable methods like ours to become even more versatile in the future.

7. CONCLUSION

We propose an expert-guided CFD meshing method that combines graph neural networks and expert heuristics to generate CFD meshes that not only improve numerical solver convergence but also capture the necessary details of flow patterns in areas with complex flow patterns. To develop our model, we introduce a novel deep-learning approach that combines the strengths of point-based models and mesh-based GNNs. Additionally, we publicly release a new annotated dataset and show how an active augmentation approach enables us to train a large deep-learning model on a small dataset without sacrificing generalization. We demonstrate that our approach outperforms the latest state-of-the-art techniques for 3D model segmentation, establishing our

model as the new state-of-the-art.

Furthermore, we introduce a mechanism for mapping predictions from a mesh to CAD surfaces and utilize conformal predictions in CAD surface classification. Our model is capable of quantifying uncertainty and handling the uncertainty with a 95% marginal statistical guarantee of correctness. Based on meshing rules established for different surface types, we utilize expert guidance to generate CFD meshes. We show that our model provides the most accurate predictions on CAD surfaces (over 99% accurate), and our conformal predictions approach leads to correct uncertainty identification and a more conservative approach to meshing with no under-refined surfaces, effectively enabling failure avoidance in the solver.

Finally, we validate the efficacy of our approach by conducting a case-study simulation of one of the aircraft in the dataset. We demonstrate that given accurate surface predictions, the resulting mesh is appropriately refined and leads to successful simulations by the solver, thus establishing the effectiveness of our approach. We make our code and data publicly available for interested readers to examine. The future work includes developing more advanced rules for meshing and expanding the dataset for a more powerful segmentation model. Additionally, efforts will be given to developing an independent CFD mesh generation pipeline for more freedom and detailed expert guidance.

ACKNOWLEDGMENTS

The authors acknowledge the Lincoln Laboratory for providing computational resources that have contributed to the research results reported in this paper. We also thank MathWorks for supporting Amin Heyrani Nobari's studies while working on this project.

REFERENCES

- [1] Chen, Wei and Ahmed, Faez. "PaDGAN: Learning to Generate High-Quality Novel Designs." *Journal of Mechanical Design* Vol. 143 No. 3 (2020). DOI [10.1115/1.4048626](https://doi.org/10.1115/1.4048626). URL https://asmedigitalcollection.asme.org/mechanicaldesign/article-pdf/143/3/031703/6632924/md_143_3_031703.pdf, URL <https://doi.org/10.1115/1.4048626>. 031703.
- [2] Chen, Wei and Ahmed, Faez. "MO-PaDGAN: Reparameterizing Engineering Designs for augmented multi-objective optimization." *Applied Soft Computing* Vol. 113 (2021): p. 107909. DOI <https://doi.org/10.1016/j.asoc.2021.107909>. URL <https://www.sciencedirect.com/science/article/pii/S1568494621008310>.
- [3] Heyrani Nobari, Amin, Chen, Wei and Ahmed, Faez. "PcDGAN: A Continuous Conditional Diverse Generative Adversarial Network For Inverse Design." *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*: pp. 606–616. 2021.
- [4] Langroudi, A Tahadjodi, Afifi, F Zare, Nobari, A Heyrani and Najafi, AF. "Modeling and numerical investigation on multi-objective design improvement of a novel cross-flow lift-based turbine for in-pipe hydro energy harvesting applications." *Energy conversion and management* Vol. 203 (2020): p. 112233.

- [5] Plewa, Tomasz, Linde, Timur and Weirs, V. Gregory. *Adaptive mesh refinement: Theory and applications: Proceedings of the chicago workshop on adaptive mesh refinement methods, September 3-5, 2003*. Springer (2005).
- [6] Fidkowski, Krzysztof J and Darmofal, David L. “Review of output-based error estimation and mesh adaptation in computational fluid dynamics.” *AIAA journal* Vol. 49 No. 4 (2011): pp. 673–694.
- [7] Becker, Roland and Rannacher, Rolf. “An optimal control approach to a posteriori error estimation in finite element methods.” *Acta Numerica* Vol. 10 (2001): p. 1–102. DOI [10.1017/S0962492901000010](https://doi.org/10.1017/S0962492901000010).
- [8] Babuška, I. and Miller, A. “A feedback finite element method with a posteriori error estimation: Part I. The finite element method and some basic properties of the a posteriori error estimator.” *Computer Methods in Applied Mechanics and Engineering* Vol. 61 No. 1 (1987): pp. 1–40. DOI [https://doi.org/10.1016/0045-7825\(87\)90114-9](https://doi.org/10.1016/0045-7825(87)90114-9). URL <https://www.sciencedirect.com/science/article/pii/S0045782587901149>.
- [9] Eriksson, Kenneth, Estep, Don, Hansbo, Peter and Johnson, Claes. “Introduction to Adaptive Methods for Differential Equations.” *Acta Numerica* Vol. 4 (1995): p. 105–158. DOI [10.1017/S0962492900002531](https://doi.org/10.1017/S0962492900002531).
- [10] Verfürth, R. “A posteriori error estimation and adaptive mesh-refinement techniques.” *Journal of Computational and Applied Mathematics* Vol. 50 No. 1 (1994): pp. 67–83. DOI [https://doi.org/10.1016/0377-0427\(94\)90290-9](https://doi.org/10.1016/0377-0427(94)90290-9). URL <https://www.sciencedirect.com/science/article/pii/S0377042794902909>.
- [11] Braack, Malte and Ern, Alexandre. “Adaptive Computation of Reactive Flows with Local Mesh Refinement and Model Adaptation.” Feistauer, Miloslav, Dolejší, Vít, Knobloch, Petr and Najzar, Karel (eds.). *Numerical Mathematics and Advanced Applications*: pp. 159–168. 2004. Springer Berlin Heidelberg, Berlin, Heidelberg.
- [12] Regenwetter, Lyle, Nobari, Amin Heyrani and Ahmed, Faez. “Deep generative models in engineering design: A review.” *Journal of Mechanical Design* Vol. 144 No. 7 (2022): p. 071704.
- [13] Huang, Keefe, Krügener, Moritz, Brown, Alistair, Menhorn, Friedrich, Bungartz, Hans-Joachim and Hartmann, Dirk. “Machine learning-based optimal mesh generation in computational fluid dynamics.” *arXiv preprint arXiv:2102.12923* (2021).
- [14] Fidkowski, Krzysztof J. and Chen, Guodong. “Metric-based, goal-oriented mesh adaptation using machine learning.” *Journal of Computational Physics* Vol. 426 (2021): p. 109957. DOI <https://doi.org/10.1016/j.jcp.2020.109957>. URL <https://www.sciencedirect.com/science/article/pii/S0021999120307312>.
- [15] Pfaff, Tobias, Fortunato, Meire, Sanchez-Gonzalez, Alvaro and Battaglia, Peter. “Learning Mesh-Based Simulation with Graph Networks.” *International Conference on Learning Representations*. 2021. URL https://openreview.net/forum?id=roNqYL0_XP.
- [16] Sadinle, Mauricio, Lei, Jing and Wasserman, Larry. “Least Ambiguous Set-Valued Classifiers With Bounded Error Levels.” *Journal of the American Statistical Association* Vol. 114 No. 525 (2019): pp. 223–234. DOI [10.1080/01621459.2017.1395341](https://doi.org/10.1080/01621459.2017.1395341). URL <https://doi.org/10.1080/01621459.2017.1395341>, URL <https://doi.org/10.1080/01621459.2017.1395341>.
- [17] Romano, Yaniv, Sesia, Matteo and Candès, Emmanuel J. “Classification with Valid and Adaptive Coverage.” *Proceedings of the 34th International Conference on Neural Information Processing Systems*. 2020. Curran Associates Inc., Red Hook, NY, USA.
- [18] Angelopoulos, Anastasios, Bates, Stephen, Malik, Jitendra and Jordan, Michael I. “Uncertainty sets for image classifiers using conformal prediction.” *arXiv preprint arXiv:2009.14193* (2020).
- [19] Su, Hang, Maji, Subhansu, Kalogerakis, Evangelos and Learned-Miller, Erik G. “Multi-view Convolutional Neural Networks for 3D Shape Recognition.” *2015 IEEE International Conference on Computer Vision (ICCV)* (2015): pp. 945–953.
- [20] Kalogerakis, Evangelos, Averkiou, Melinos, Maji, Subhansu and Chaudhuri, Siddhartha. “3D shape segmentation with projective convolutional networks.” *proceedings of the IEEE conference on computer vision and pattern recognition*: pp. 3779–3788. 2017.
- [21] Wu, Zhirong, Song, Shuran, Khosla, Aditya, Yu, Fisher, Zhang, Linguang, Tang, Xiaoou and Xiao, Jianxiong. “3D ShapeNets: A deep representation for volumetric shapes.” *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*: pp. 1912–1920. 2015. DOI [10.1109/CVPR.2015.7298801](https://doi.org/10.1109/CVPR.2015.7298801).
- [22] Brock, Andrew, Lim, Theodore, Ritchie, James M. and Weston, Nick. “Generative and Discriminative Voxel Modeling with Convolutional Neural Networks.” *ArXiv* Vol. abs/1608.04236 (2016).
- [23] Tchapmi, Lyne P., Choy, Christopher Bongsoo, Armeni, Iro, Gwak, JunYoung and Savarese, Silvio. “SEGCloud: Semantic Segmentation of 3D Point Clouds.” *2017 International Conference on 3D Vision (3DV)* (2017): pp. 537–547.
- [24] Hanocka, Rana, Fish, Noa, Wang, Zhenhua, Giryas, Raja, Fleishman, Shachar and Cohen-Or, Daniel. “ALIGNet: Partial-Shape Agnostic Alignment via Unsupervised Learning.” *ACM Trans. Graph.* Vol. 38 (2018): pp. 1:1–1:14.
- [25] Maturana, Daniel and Scherer, Sebastian. “VoxNet: A 3D Convolutional Neural Network for real-time object recognition.” *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*: pp. 922–928. 2015. DOI [10.1109/IROS.2015.7353481](https://doi.org/10.1109/IROS.2015.7353481).
- [26] Sun, Kai, Zhang, Jianshe, Liu, Junmin, Yu, Ruixuan and Song, Zengjie. “DRCNN: Dynamic Routing Convolutional Neural Network for Multi-View 3D Object Recognition.” *Trans. Img. Proc.* Vol. 30 (2021): p. 868–877. DOI [10.1109/TIP.2020.3039378](https://doi.org/10.1109/TIP.2020.3039378). URL <https://doi.org/10.1109/TIP.2020.3039378>.

- [27] Qi, Charles R., Su, Hao, Mo, Kaichun and Guibas, Leonidas J. “PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation.” *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2017.
- [28] Qi, Charles Ruizhongtai, Yi, Li, Su, Hao and Guibas, Leonidas J. “PointNet++: Deep Hierarchical Feature Learning on Point Sets in a Metric Space.” Guyon, I., Luxburg, U. Von, Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S. and Garnett, R. (eds.). *Advances in Neural Information Processing Systems*, Vol. 30. 2017. Curran Associates, Inc. URL <https://proceedings.neurips.cc/paper/2017/file/d8bf84be3800d12f74d8b05e9b89836f-Paper.pdf>.
- [29] Wang, Yue, Sun, Yongbin, Liu, Ziwei, Sarma, Sanjay E., Bronstein, Michael M. and Solomon, Justin M. “Dynamic Graph CNN for Learning on Point Clouds.” *ACM Transactions on Graphics (TOG)* (2019).
- [30] Huang, Qiangui, Wang, Weiyue and Neumann, Ulrich. “Recurrent slice networks for 3d segmentation of point clouds.” *Proceedings of the IEEE conference on computer vision and pattern recognition*: pp. 2626–2635. 2018.
- [31] Li, Yangyan, Bu, Rui, Sun, Mingchao, Wu, Wei, Di, Xinhan and Chen, Baoquan. “PointCNN: Convolution On X-Transformed Points.” Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N. and Garnett, R. (eds.). *Advances in Neural Information Processing Systems*, Vol. 31. 2018. Curran Associates, Inc. URL <https://proceedings.neurips.cc/paper/2018/file/f5f8590cd58a54e94377e6ae2eded4d9-Paper.pdf>.
- [32] Wu, Wenxuan, Qi, Zhongang and Fuxin, Li. “PointConv: Deep Convolutional Networks on 3D Point Clouds.” *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 2019.
- [33] Ma, Xu, Qin, Can, You, Haoxuan, Ran, Haoxi and Fu, Yun. “Rethinking Network Design and Local Geometry in Point Cloud: A Simple Residual MLP Framework.” *International Conference on Learning Representations*. 2022. URL https://openreview.net/forum?id=3Pbra-_u76D.
- [34] Hanocka, Rana, Hertz, Amir, Fish, Noa, Giryas, Raja, Fleishman, Shachar and Cohen-Or, Daniel. “MeshCNN: a network with an edge.” *ACM Transactions on Graphics (TOG)* Vol. 38 (2019): pp. 1 – 12.
- [35] Lian, Chunfeng, Wang, Li, Wu, Tai-Hsien, Wang, Fan, Yap, Pew-Thian, Ko, Ching-Chang and Shen, Dinggang. “Deep Multi-Scale Mesh Feature Learning for Automated Labeling of Raw Dental Surfaces From 3D Intraoral Scanners.” *IEEE Transactions on Medical Imaging* Vol. 39 No. 7 (2020): pp. 2440–2450. DOI [10.1109/TMI.2020.2971730](https://doi.org/10.1109/TMI.2020.2971730).
- [36] Zheng, Youyi, Chen, Beijia, Shen, Yuefan and Shen, Kaidi. “TeethGNN: Semantic 3D Teeth Segmentation with Graph Neural Networks.” *IEEE Transactions on Visualization and Computer Graphics* (2022): pp. 1–1 DOI [10.1109/TVCG.2022.3153501](https://doi.org/10.1109/TVCG.2022.3153501).
- [37] Wen, Tingxi, Zhuang, Jiafu, Du, Yu, Yang, Linjie and Xu, Jianfei. “Dual-Sampling Attention Pooling for Graph Neural Networks on 3D Mesh.” *Computer Methods and Programs in Biomedicine* Vol. 208 (2021): p. 106250. DOI <https://doi.org/10.1016/j.cmpb.2021.106250>. URL <https://www.sciencedirect.com/science/article/pii/S0169260721003242>.
- [38] Qi, Xiaojuan, Liao, Renjie, Jia, Jiaya, Fidler, Sanja and Urtasun, Raquel. “3D Graph Neural Networks for RGBD Semantic Segmentation.” *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*. 2017.
- [39] Karras, Tero, Aittala, Miika, Hellsten, Janne, Laine, Samuli, Lehtinen, Jaakko and Aila, Timo. “Training Generative Adversarial Networks with Limited Data.” (2020). DOI [10.48550/ARXIV.2006.06676](https://doi.org/10.48550/ARXIV.2006.06676). URL <https://arxiv.org/abs/2006.06676>.
- [40] Nobari, Amin Heyrani, Chen, Wei and Ahmed, Faez. “Range-Constrained Generative Adversarial Network: Design Synthesis Under Constraints Using Conditional Generative Adversarial Networks.” *Journal of Mechanical Design* Vol. 144 No. 2 (2021). DOI [10.1115/1.4052442](https://doi.org/10.1115/1.4052442). URL https://asmedigitalcollection.asme.org/mechanicaldesign/article-pdf/144/2/021708/6768330/md_144_2_021708.pdf, URL <https://doi.org/10.1115/1.4052442>. 021708.
- [41] Veličković, Petar, Cucurull, Guillem, Casanova, Arantxa, Romero, Adriana, Lio, Pietro and Bengio, Yoshua. “Graph attention networks.” *arXiv preprint arXiv:1710.10903* (2017).
- [42] McDonald, Robert A. and Gloude-mans, James R. *Open Vehicle Sketch Pad: An Open Source Parametric Geometry and Analysis Tool for Conceptual Aircraft Design*: DOI [10.2514/6.2022-0004](https://doi.org/10.2514/6.2022-0004). URL <https://arc.aiaa.org/doi/pdf/10.2514/6.2022-0004>, URL <https://arc.aiaa.org/doi/abs/10.2514/6.2022-0004>.
- [43] Cignoni, Paolo, Callieri, Marco, Corsini, Massimiliano, Dellepiane, Matteo, Ganovelli, Fabio and Ranzuglia, Guido. “MeshLab: an Open-Source Mesh Processing Tool.” Scarano, Vittorio, Chiara, Rosario De and Erra, Ugo (eds.). *Eurographics Italian Chapter Conference*. 2008. The Eurographics Association. DOI [10.2312/LocalChapterEvents/ItalChap/ItalianChapConf2008/129-136](https://doi.org/10.2312/LocalChapterEvents/ItalChap/ItalianChapConf2008/129-136).
- [44] Hoppe, Hugues, DeRose, Tony, Duchamp, Tom, McDonald, John and Stuetzle, Werner. “Mesh optimization.” *Proceedings of the 20th annual conference on Computer graphics and interactive techniques*: pp. 19–26. 1993.
- [45] Kipf, Thomas N and Welling, Max. “Semi-supervised classification with graph convolutional networks.” *arXiv preprint arXiv:1609.02907* (2016).

APPENDIX A. EXPERT-GUIDED RULES FOR DIFFERENT AIRCRAFT SURFACES

In this section, we present the specific rules provided by experts we use for the flight conditions and software mentioned in section 3.3. Below are the specific mesh settings that are applied in PointWise meshing software for different aircraft parts:

1. Wings:

- Set surface mesh dimension = 0.05

- Set surface mesh to: quadrilateral cells dominant
- Set volume mesh to: hexahedral cells dominant
- Set initial wall spacing = $4.7e-6$
- Set Growth rate = 1.1
- Set Collision Buffer = 2.0
- Grow volume cells from surface cells. Check Y+ value... If > 1 , then decrease initial wall spacing proportionally

2. Stabalizers:

- Set surface mesh dimension = 0.2
- Set surface mesh to: quadrilateral cells dominant
- Set volume mesh to: hexahedral cells dominant
- Set initial wall spacing = $4.7e-6$
- Set Growth rate = 1.1

- Set Collision Buffer = 2.0
- Grow volume cells from surface cells. Check Y+ value... If > 1 , then decrease initial wall spacing proportionally

3. Fuselage:

- Set surface mesh dimension = 1.0
- Set surface mesh to: quadrilateral cells dominant
- Set volume mesh to: hexahedral cells dominant
- Set initial wall spacing = $4.7e-6$
- Set Growth rate = 1.1
- Set Collision Buffer = 2.0
- Grow volume cells from surface cells. Check Y+ value... If > 1 , then decrease initial wall spacing proportionally