BTP Report

Shafi Ur Rahman Khan

CS21B061

Guide: Dr. Chalavadi Vishnu

Indian Institute of Technology, Tirupati

Declaration

I declare that this written submission represents my ideas in my own words and where others' ideas or words have been included, I have adequately cited and referenced the original sources. I also declare that I have adhered to all principles of academic honesty and integrity and have not misrepresented or fabricated or falsified any idea/data/fact/source in my submission to the best of my knowledge. I understand that any violation of the above will be cause for disciplinary action by the Institute and can also evoke penal action from the sources which have thus not been properly cited or from whom proper permission has not been taken when needed.

Place: Tirupati
Date: 08-05-2025

Signature
Shafi Ur Rahman Khan
CS21B061

Bona Fide Certificate

This is to certify that the thesis titled **POINTNET, 3DETR, CONVERSION OF MESH TO POINT CLOUDS AND VICE VERSA**, submitted by **Shafi ur Rahman Khan**, to the Indian Institute of Technology, Tirupati, for the award of the degree of , is a bona fide record of the work done by them under my supervision. The contents of this , in full or in parts, have not been submitted to any other Institute or University for the award of any degree or diploma.

C. Vishnu

Place: Tirupati
Date: 08-05-2025

Dr. Chalavadi Vishnu
Guide
Assistant Professor
Department of Computer Science of Engineering
IIT Tirupati - 517619

Abstract

This project presents a comprehensive study and practical implementation of a bidirectional conversion pipeline between 3D mesh files and point cloud representations. The motivation stems from the observation that many classical geometry-based algorithms operate on mesh models, while modern deep learning approaches and sensor data pipelines favor point clouds due to their simplicity and processing efficiency. This duality requires robust methods for transforming data between these two formats to support both algorithmic compatibility and real-world applications.

The project begins by exploring deep learning models tailored for point clouds. PointNet, a classification model designed for unordered point sets, was trained on the ModelNet10 dataset and achieved a training accuracy of 87.63% and a validation accuracy of 85.87%, demonstrating its effectiveness in geometric recognition tasks. Further, 3DETR (3D Detection Transformer) was studied as an end-to-end object detection framework that leverages self-attention mechanisms to learn contextual relationships within 3D point sets, offering richer feature representations than traditional MLP-based architectures.

Following the model analysis, a mesh sampling pipeline was developed to convert 3D triangular meshes into point clouds. The sampling process ensures an even distribution of points across mesh surfaces based on triangle area, with optional voxelization to reduce redundancy and normalize data for learning tasks. The sampled point clouds were then subjected to surface reconstruction using two classical methods: Poisson Surface Reconstruction (PSR) and Alpha Shapes. PSR employs a global implicit function derived from the Poisson equation to produce watertight surfaces, while Alpha Shapes use a localized, parameter-controlled approach to capture sharper features and fine details.

To assess reconstruction quality, we employed three evaluation metrics: Chamfer Distance, Hausdorff Distance, and Surface Area Error. Experimental results showed that Alpha Shapes with $\alpha = 10$ and $\alpha = 20$ provided lower Chamfer and Hausdorff distances, indicating closer geometric alignment with the original meshes. In contrast, Poisson reconstruction exhibited a lower surface area error, suggesting better global surface closure but higher deviation in local geometry.

This work establishes a flexible and effective pipeline for converting between mesh and point cloud formats. It enables compatibility with a broad range of 3D algorithms, supports machine learning and geometry-based workflows, and provides empirical insights into the trade-offs involved in surface reconstruction techniques.

Contents

Point Clouds and Mesh Files

June 6, 2025

1 Previous BTP Work

1.1 PointNet: A Classification Model for Point Clouds

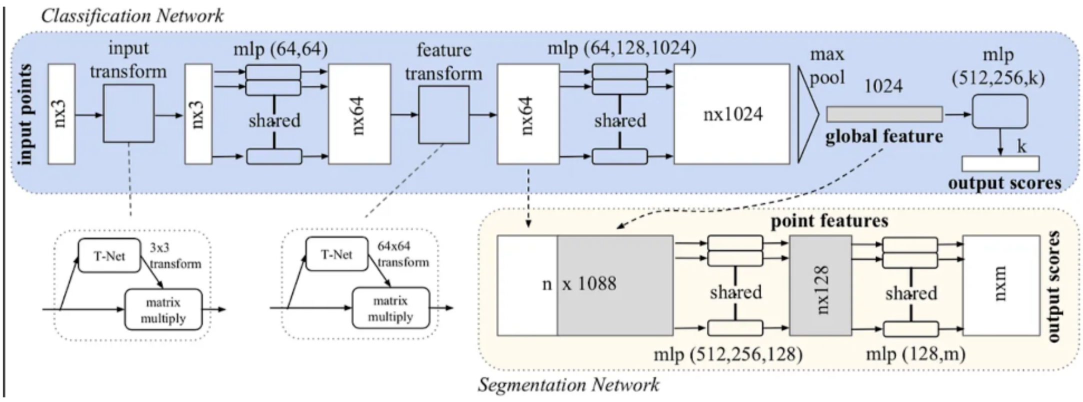

Figure 1: Architecture of the model

1.1.1 Results:

The model gave the **training accuracy of 87.63%** and **validation accuracy of 85.87%**. Sparse categorical accuracy metric was used to measure classification performance. The model was trained on the **ModelNet10** dataset and demonstrated strong classification performance.

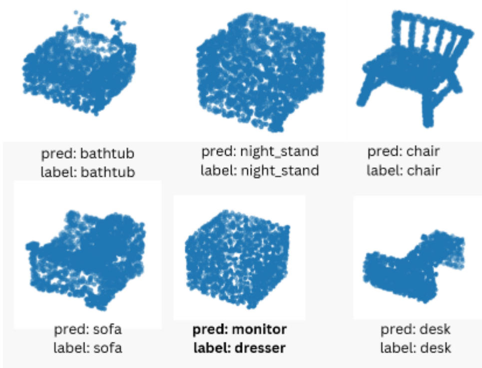

Figure 2: Results

1.2 3DETR: Transformer-Based Object Detection in 3D Point Clouds

3DETR (3D Detection Transformer) is an end-to-end object detection model leveraging transformer networks for 3D point clouds. Unlike traditional MLP-based architectures such as PointNet, 3DETR employs self-attention mechanisms to capture complex spatial relationships.

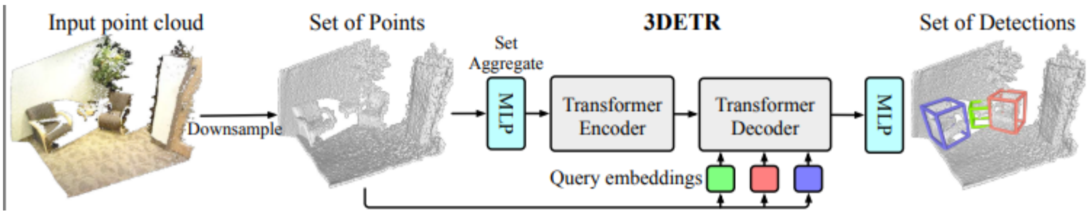

Figure 3: 3D Detection Transformer

1.2.1 Model Components:

The **Transformer encoder** processes individual points and learns feature embeddings.The **Transformer decoder** generates query-based predictions for the bounding boxes. Simultaneously, **Hungarian loss function** ensures one-to-one mapping between predicted and ground-truth objects.

Features	PointNet	3DETR
Feature Extraction	MLP-Based	Transformer-Based
Positional Encoding	XYZ coordinates	Fourier embeddings
Decoder	None	Transformer decoder
Loss Function	Task-specific losses	Set-based loss (Hungarian)

Figure 4: DIfferences between 3DETR and PointNet

The comparison table in the image highlights how 3DETR differs from PointNet. Although PointNet uses MLP-based feature extraction, 3DETR employs Fourier embeddings and transformer-based decoding for richer contextual understanding.

2 Mesh Sampling

2.1 Overview

Mesh sampling converts a 3D triangular mesh into a point cloud for applications such as 3D modeling. The goal is a uniform distribution of points across the mesh surface.

2.2 Steps in Mesh Sampling

1. **Mesh Loading:** Load the 3D triangular mesh.
2. **Surface Area Calculation:** Compute the surface area of each triangle.
3. **Point Distribution:** Distribute points based on triangle surface area using uniform/random sampling.
4. **Point Cloud Generation:** Collect distributed points to form the final point cloud.

2.3 Key Considerations

1. **Uniformity:** Ensure even point distribution to avoid clustering.
2. **Density:** Adjust density based on application needs. High density for detailed scans, low density for visualization.
3. **Computational Efficiency:** Optimize for large meshes without excessive computational overhead.

2.4 Practical Implementation

The ModelNet40 dataset is used as input, successfully converting the mesh data into point clouds.

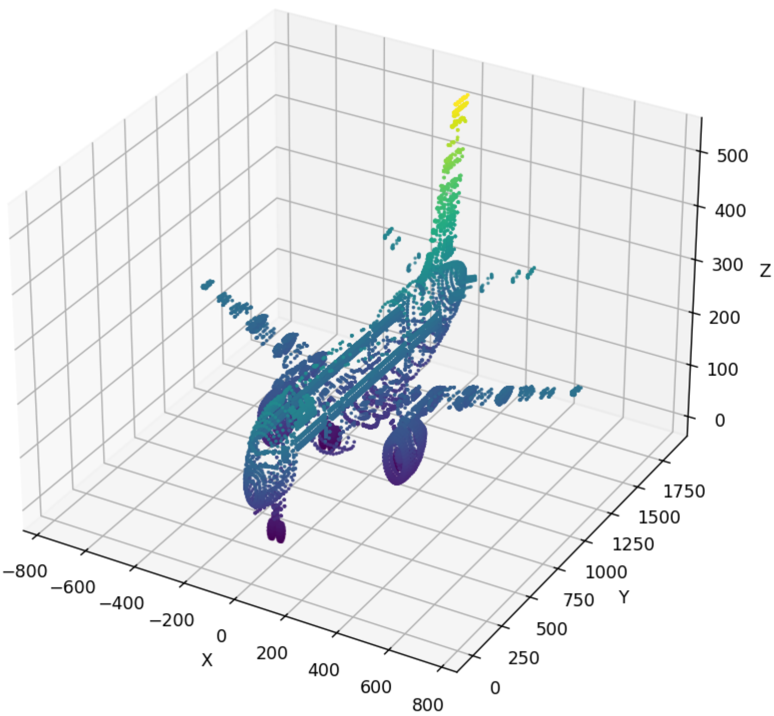

Figure 5: Input mesh from ModelNet40

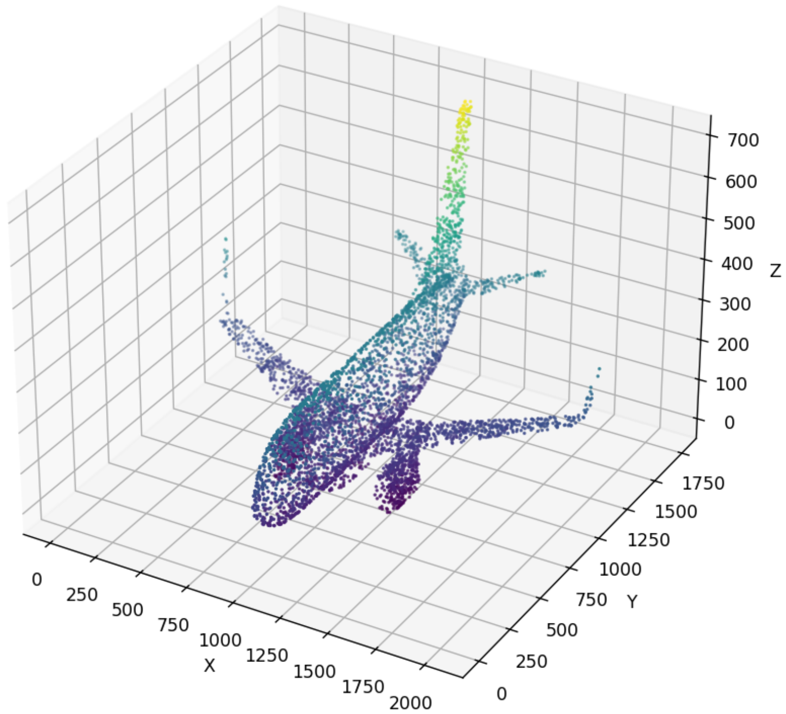

Figure 6: Point Cloud Output

3 Surface Reconstruction

3.1 Overview

Surface reconstruction converts point clouds into continuous surfaces, which is essential in computer graphics, robotics, and 3D modeling. Methods include Poisson Surface Reconstruction and Alpha Shapes Reconstruction.

3.2 Steps in Surface Reconstruction

- Load point cloud data.
- Estimate the point normals for orientation.
- **Poisson Surface Reconstruction:** Uses Poisson equation for smooth surface generation.
- **Alpha Shapes Reconstruction:** Uses alpha shapes for sharper edges and features.
- Visualize results to evaluate reconstruction quality.

3.2.1 Poisson Surface Reconstruction

Poisson Surface Reconstruction is a technique used to convert a point cloud with normals to a surface mesh. It is particularly effective in applications such as 3D scanning, computer graphics, and reverse engineering.

Theory The core idea of PSR is to treat the input point cloud as samples of the gradient of an implicit function χ representing the surface. The goal is to reconstruct χ by solving a Poisson equation that best matches the gradient field formed from the normals.

$$\Delta\chi = \nabla \cdot \mathbf{V}$$

Here,

- χ is the implicit surface function whose zero level set defines the reconstructed surface.
- $\mathbf{V}$ is the vector field derived from the input point normals.
- $\nabla \cdot \mathbf{V}$ is the divergence of the vector field.

Solving this equation yields an implicit representation of the surface. A mesh is then extracted from this representation.

Limitations

- Requires accurate **normal estimation**; poor normals degrade reconstruction quality.
- Assumes a fairly **uniform sampling density**.
- May **oversmooth** fine details or sharp edges.

Practical Use in The Work In this project, PSR was implemented using Open3D with the following call:

```
poisson_mesh, densities =

o3d.geometry.TriangleMesh.create_from_point_cloud_poisson(pcd, depth=9)
```

The reconstruction was applied to voxelized point clouds generated from ModelNet40 meshes. The evaluation revealed higher Chamfer and Hausdorff distances compared to alpha-shape methods, indicating geometric mismatches likely caused by over-smoothing and normal inaccuracies.

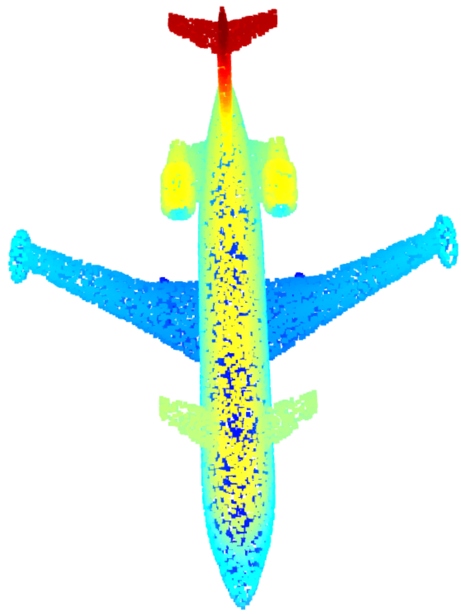

Figure 7: Point Cloud Input

Figure 8: PSR Output (From a Different Angle)

3.2.2 Alpha Shapes Reconstruction

Alpha Shapes is a geometric surface reconstruction algorithm that extends the concept of a convex hull to allow reconstruction of non-convex and detailed surfaces from point clouds. Convex hull is the smallest convex polygon (or polyhedron in higher dimensions) that encloses a set of points. In the context of alpha reconstruction, it's often used as a baseline for constructing more intricate shapes. It is especially useful when working with sparse, noisy, or irregularly distributed data.

Theory The algorithm constructs a Delaunay triangulation from the point cloud and then filters simplices (edges, triangles, and tetrahedra) based on a user-defined parameter α. Only those simplices whose circumscribed sphere has a radius less than or equal to $1/\alpha$ are retained. The remaining triangles define the reconstructed surface.

Delaunay triangulation is a method for creating a triangulation of a set of points, where the triangles' circumcircles don't contain any other points of the set.

Parameter (α) The value of α significantly influences the reconstruction:

- **Small** α values preserve sharp features but may lead to overfitting or disconnected components.
- **Large** α values produce smoother and more connected surfaces but may miss finer details.

Practical Implementation In our work, Open3D's Alpha Shape reconstruction function was used:

```
alpha_mesh =

o3d.geometry.TriangleMesh.create_from_point_cloud_alpha_shape(pcd, alpha=10)
```

Meshes were generated using multiple α values (5, 10, 20) to observe their effect on reconstruction quality.

Evaluation Based on Chamfer Distance, Hausdorff Distance, and Surface Area Error metrics, we found that:

- $\alpha = 10$ provided the best Chamfer Distance (22.55), indicating good average accuracy.
- $\alpha = 20$ yielded the best Hausdorff Distance (21.25) and Surface Area similarity, suggesting better overall geometric consistency.
- $\alpha = 5$ resulted in poor performance across all metrics, likely due to overfitting and excessive surface detail.

Limitations While Alpha Shapes preserve sharp geometry, they do not guarantee watertight meshes. This was reflected in the inability to compute volume metrics in our evaluation due to non-watertight outputs.

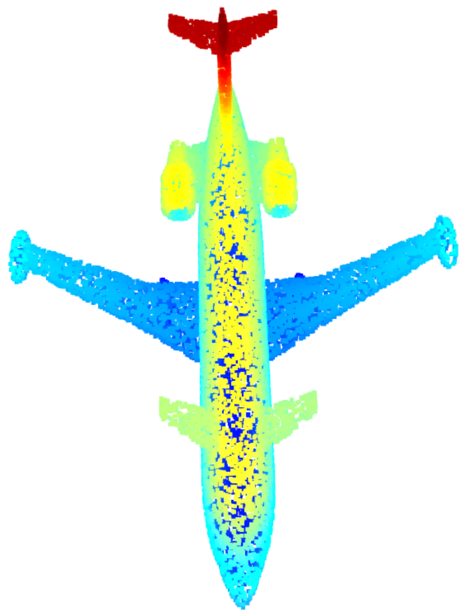

Figure 9: Point Cloud Input

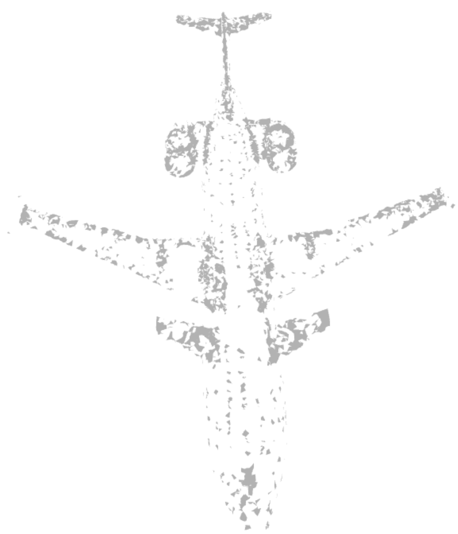

Figure 10: Alpha reconstruction output(alpha = 5)

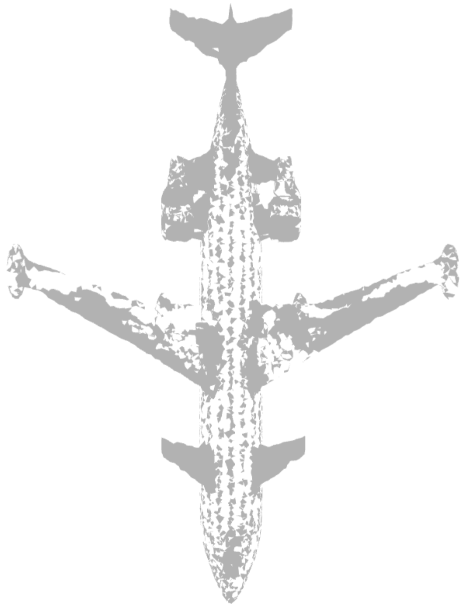

Figure 11: Alpha reconstruction output(alpha = 10)

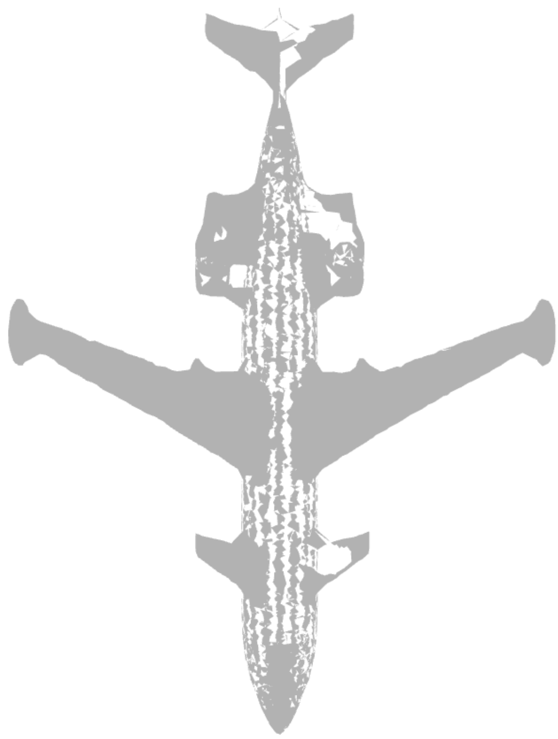

Figure 12: Alpha reconstruction output(alpha = 20)

4 Evaluation Metrics

To quantify the quality of reconstructed surfaces, we employ the following standard geometric evaluation metrics:

4.1 Chamfer Distance (CD)

Chamfer Distance measures the average bidirectional closeness between two point clouds. Given two point sets P and Q, the Chamfer Distance is defined as:

$$CD(P, Q) = \frac{1}{|P|} \sum_{p \in P} \min_{q \in Q} \|p - q\|_2 + \frac{1}{|Q|} \sum_{q \in Q} \min_{p \in P} \|q - p\|_2$$

A lower Chamfer Distance indicates better alignment between the reconstructed and original surfaces. It is robust to noise and missing points.

4.2 Hausdorff Distance (HD)

Hausdorff Distance measures the maximum discrepancy between two point clouds. It captures the largest minimum distance from any point in one set to the other:

$$HD(P, Q) = \max \left\{ \sup_{p \in P} \inf_{q \in Q} \|p - q\|_2, \sup_{q \in Q} \inf_{p \in P} \|q - p\|_2 \right\}$$

It reflects the worst-case error between the two surfaces. Lower HD values indicate tighter alignment with minimal outlier deviations.

4.3 Surface Area Error

Surface Area Error measures the absolute difference in surface areas between the original and reconstructed meshes:

$$\text{Area Error} = |A_{original} - A_{reconstructed}|$$

A low surface area error indicates that the reconstruction preserves the geometric scale and proportions of the original shape.

4.4 Implementation

All metrics were implemented using Open3D and SciPy libraries. Point clouds were uniformly sampled from both the original and reconstructed meshes before evaluation.

4.5 Visualization

To better understand local reconstruction accuracy, we used heatmaps to visualize point-wise distances between the original and reconstructed point clouds. Color maps (e.g., `jet`) were applied to highlight high-error regions.

5 Results and Inferences

Method	Chamfer Distance	Hausdorff Distance	Surface Area Error
Poisson	**34,228.84**	**34.79**	**614,466.29**
Alpha (α=5)	156.63	**54.89**	1,205,951.71
Alpha (α=10)	**22.55**	21.80	814,950.36
Alpha (α=20)	25.59	**21.25**	**687,640.06**

Figure 13: Evaluation Table with Different Metrics

- Poisson reconstruction, while producing geometrically smooth and visually appealing meshes, tends to overfit or incorrectly interpolate parts of the shape — especially for sparse or noisy point clouds — leading to very high Chamfer Distance. Its lower surface area error suggests it's closing the shape better.
- Alpha Shapes ($\alpha = 10$ or $\alpha = 20$) outperform Poisson quantitatively in Chamfer and Hausdorff metrics, implying better local accuracy to the original shape. However, they struggle more in maintaining correct surface area, possibly due to underfitting or noisy boundary surfaces.
- Alpha ($\alpha = 5$) overfits the points severely and introduces a large number of tiny triangles and noise, leading to high error in both geometry and surface area.
- $\alpha = 10$ provided the best Chamfer Distance (22.55), indicating good average accuracy.
- $\alpha = 20$ yielded the best Hausdorff Distance (21.25) and Surface Area similarity, suggesting better overall geometric consistency.
- $\alpha = 5$ resulted in poor performance across all metrics, likely due to overfitting and excessive surface detail.

6 Future Work

1. Implement selective conversion within bounding boxes for targeted processing.

7 References

1. Matthew Berger, Andrea Tagliasacchi, Lee Seversky, Pierre Alliez, Joshua Levine, et al.. State of the Art in Surface Reconstruction from Point Clouds. Eurographics 2014 - State of the Art Reports, Apr 2014, Strasbourg, France. pp.161-185, ff10.2312/egst.20141040ff. ffhal-01017700f

2. Kazhdan, M., Bolitho, M. and Hoppe, H., 2006, June. Poisson surface reconstruction. In Proceedings of the fourth Eurographics symposium on Geometry processing (Vol. 7, No. 4).