

Multi-Modal 3D Mesh Reconstruction from Images and Text

Melvin Reka¹, Tessa Pulli¹, and Markus Vincze¹

Abstract—6D object pose estimation for unseen objects is essential in robotics but traditionally relies on trained models that require large datasets, high computational costs, and struggle to generalize. Zero-shot approaches eliminate the need for training but depend on pre-existing 3D object models, which are often impractical to obtain. To address this, we propose a language-guided few-shot 3D reconstruction method, reconstructing a 3D mesh from few input images. In the proposed pipeline, receives a set of input images and a language query. A combination of GroundingDINO and Segment Anything Model outputs segmented masks from which a sparse point cloud is reconstructed with VGGSfM. Subsequently, the mesh is reconstructed with the Gaussian Splatting method SuGAR. In a final cleaning step, artifacts are removed, resulting in the final 3D mesh of the queried object. We evaluate the method in terms of accuracy and quality of the geometry and texture. Furthermore, we study the impact of imaging conditions such as viewing angle, number of input images, and image overlap on 3D object reconstruction quality, efficiency, and computational scalability.

Index Terms—Vision Language Models, Language-guided Reconstruction, Few-shot Reconstruction

I. INTRODUCTION

6D object pose estimation for unseen objects is a critical task in robotics. Traditional methods estimate instance object poses using trained networks [1–4]. However, training models for object pose estimation is a limitation as it requires large annotated datasets, has high computational costs, and encounters difficulties in generalizing to unknown objects or environments [5]. An alternative to these methods are training-free zero-shot approaches [6, 7]. Methods such as ZS6D [6] utilize a ground-truth object model to find 2D-3D correspondences between the model and the images from

which the pose is computed using a PnP algorithm [8]. Although these methods offer a considerable advantage due to their zero-shot capabilities, they are limited by the requirement for a ground-truth 3D object model [2, 6]. Obtaining high-quality 3D models can be labor-intensive, expensive, and impractical for large-scale or real-time applications [9]. With the advent of diffusion models, recent works have proposed methods for few-shot [10, 11] or even single-shot [10, 12] 3D model reconstruction based on images. By combining SuGAR [11] with Segment Anything Model [13] (SAM) and GroundingDINO [14], we introduce a novel method to reconstruct 3D models based on images and language prompts. As input, we receive several RGB images of a scene. According to the language input, the queried object is detected with GroundingDINO [14]. Based on the bounding box, SAM [13] generates masked images depicting only the queried object. This serves as input for a sparse reconstruction with VGGSfM [15]. The sparse point cloud is then processed by SuGAR [11] to reconstruct a 3D mesh. In an automated post-processing step, artefacts are removed, resulting in a cleaned mesh. Finally, the reconstruction is evaluated on several experiments in terms of accuracy and quality of the geometric reconstruction as well as the reconstructed texture.

In summary, the paper has the following key contributions:

- We propose a novel language-guided few-shot reconstruction approach that allows 3D model reconstruction.
- Evaluation of the few-shot reconstruction method, including an analysis of the required input images considering efficiency and performance of the reconstruction.

The rest of this work is organized as follows: Section II introduces the related work of 6D object pose estimation, 3D model reconstruction, and language-guided segmentation. Section III describes the pipeline of our proposed methods. In

¹ all authors are with the Automation and Control Institute, TU Wien Vienna, Austria: e12102393@student.tuwien.ac.at; {pulli, vincze}@acin.ac.tuwien.at




Fig. 1. Images of an object, accompanied by a descriptive text prompt, are processed through the pipeline. A sparse reconstruction using COLMAP via VGGSfM is then performed to generate and refine the mesh with SuGaR.

Section IV, the experimental setup is presented while Section V discusses the evaluation. Section VI concludes the paper with a summary and outlook.

II. RELATED WORK

In this section, we discuss the related work by revisiting 6D object pose estimation, few-shot reconstruction methods, and language-guided object segmentation.

A. 6D Object Pose Estimation

6D object pose estimation of unseen objects is a core task in robotics. Classical methods estimate the pose of objects using trained networks for object instances [1–3]. These methods require large annotated datasets, which are costly to acquire while requiring high computational power during training [5]. Furthermore, these models struggle to generalize to unseen objects [1, 3, 4], limiting their applicability in real-world scenarios. These challenges can be overcome with zero-shot object pose estimation methods [6, 7]. These approaches eliminate the need for extensive training by leveraging prior knowledge by assuming that object models exist. Given the reference model, 2D-3D correspondences between the object’s model and a set of input images can be established [8]. With this information available, the object pose is computed by a PnP/RANSAC algorithm [16]. Although zero-shot methods offer a significant advantage by overcoming the training phase, they still require high-quality ground-truth 3D models. This dependency presents challenges in practical applications, as the acquisition of accurate 3D models is labor-intensive and requires expensive equipment [17].

B. Few-Shot 3D Reconstruction Methods

Recent advances in generative models, particularly diffusion models, have opened new possibilities to acquire 3D meshes. Works such as Wonder3D [12], Gaussian Splatting [18], Dreamfusion [19], and Sugar [11] have demonstrated the potential of generating 3D object representations from a limited number of 2D images. Typically, these methods use 2D diffusion models to generate novel views from different camera poses. From these novel views, a 3D model is reconstructed with a stochastic 3D reconstruction framework [19] or can be parameterized as a voxel radiance field from which the mesh is extracted with a marching cubes procedure [20]. Because existing reconstruction methods typically assume clean, isolated object inputs, several methods introduced a pre-processing step to delete the background of the input images to avoid noise while reconstructing the images [21].

C. Language Guided Object Segmentation

The integration of vision-language models (VLMs) such as CLIP [22] has significantly expanded the capabilities of computer vision systems, allowing them to understand and process images based on textual descriptions. CLIP [22] has been incorporated into a wide range of applications, allowing methods including scene understanding [23], object

recognition [24], and generative modeling [19]. An area where this integration has proven particularly beneficial is object segmentation [13], where language-guided approaches allow for more intuitive and adaptable object selection. Recent advances in segmentation models provide a method to segment objects in an image using minimal user input. SAM [13] allows for object selection through different means, such as points, bounding boxes, or language-based prompts, making it an effective tool for isolating objects in complex scenes. By leveraging SAM’s capabilities, objects can be accurately segmented and masked before being passed into a 3D reconstruction pipeline.

III. METHOD

As shown in Figure 1, images of the object, along with a descriptive text prompt, are processed through the pipeline. LangSAM [25] combines GroundingDINO [14] and SAM [13] for text-driven segmentation. GroundingDINO processes the text prompt to generate bounding boxes around relevant objects, which SAM then uses to create binary masks with a 50-pixel padding. This padding proved helpful when handling semi-transparent objects and reduces artifacts introduced during mesh reconstruction. Focusing on essential scan areas enhances reconstruction quality while maintaining computational efficiency.

The generated masks, along with the original images, are then used for 3D reconstruction with VGGSfM [15]. As a fully differentiable structure-from-motion pipeline, VGGSfM estimates camera parameters, determines camera positions, and reconstructs a sparse point cloud by tracking corresponding 2D points across multiple views. This end-to-end differentiable approach enhances the accuracy and robustness of the reconstruction process by eliminating the need to chain pairwise matches and enabling simultaneous recovery these[15].

Using the resulting COLMAP[26] dataset and the extracted RGB masks, a textured mesh is generated with SuGaR[11], which employs Gaussian Splatting to efficiently optimize and extract a high-resolution 3D surface. However, since SuGaR can introduce artifacts during mesh generation, an automated script utilizing PyMeshLab[27] is applied to remove these artifacts.

IV. EXPERIMENTS

In this section, the experimental setup is discussed. During our evaluation, we assess the accuracy and quality of the geometric construction as well as the reconstructed texture.

A. Implementation details:

All experiments are conducted on a system equipped with an AMD Ryzen 9 5950X CPU, 128GB RAM, and an NVIDIA RTX 3090 GPU with 24GB VRAM. The implementation is containerized using Docker to ensure reproducibility across different hardware environments.

B. Evaluation Metrics

To assess the accuracy and quality of both geometric reconstruction and texture extraction, we distinguish between geometric metrics and texture similarity.

Geometric Metrics

We evaluate the reconstructed 3D geometry using Chamfer Distance [28] and Intersection over Union [29]. CD quantifies the average squared distance between nearest neighbors in the predicted and ground-truth meshes. It is defined as:

$$CD(P, Q) = \frac{1}{|P|} \sum_{p \in P} \min_{q \in Q} \|p - q\|^2 + \frac{1}{|Q|} \sum_{q \in Q} \min_{p \in P} \|q - p\|^2 \quad (1)$$

where P and Q are the sets of points in the predicted and ground-truth meshes. Low CD values indicate a more precise geometric alignment.

IoU measures the volumetric similarity between the reconstructed and ground-truth models:

$$IoU = \frac{|V_P \cap V_Q|}{|V_P \cup V_Q|} \quad (2)$$

where V_P and V_Q represent the volumetric reconstructions of the predicted and ground-truth meshes. IoU measures the proportion of the shared volume, with higher values indicating better alignment.

Texture Similarity

To assess texture extraction accuracy, we use the three key metrics employed in SuGaR [11].

The Peak Signal-to-Noise Ratio is defined as:

$$PSNR = 10 \log_{10} \left(\frac{MAX^2}{MSE} \right) \quad (3)$$

where MAX is the maximum possible pixel value, and MSE is the mean squared error between the predicted and ground-truth textures. Higher values indicate better pixel-wise preservation, but do not reflect human perception.

The Structural Similarity Index [30] takes into account luminance, contrast, and texture integrity, reflecting human perception. SSIM is computed as:

$$SSIM(x, y) = \frac{(2\mu_x\mu_y + C_1)(2\sigma_{xy} + C_2)}{(\mu_x^2 + \mu_y^2 + C_1)(\sigma_x^2 + \sigma_y^2 + C_2)} \quad (4)$$

where μ_x and μ_y are the mean intensities, σ_x^2 and σ_y^2 are the variances, and σ_{xy} is the covariance between the predicted and ground-truth images.

The Learned Perceptual Image Patch Similarity [31] is given by:

$$LPIPS(x, y) = \sum_l w_l \|F_l(x) - F_l(y)\|_2^2 \quad (5)$$

where F_l represents the feature maps at layer l of a pretrained network, and w_l are learned weights. LPIPS captures high-level perceptual differences, making it effective for identifying distortions and artifacts beyond pixel-wise comparisons.

C. Experimental Setup

Our experiments evaluate the impact of various imaging conditions on the quality of 3D object reconstruction. We investigate how the viewing angle θ influences feature extraction and reconstruction accuracy, as different angles affect feature visibility. We also examine the effect of the number of input images on reconstruction convergence, assessing how multi-view stereo improves model quality. The overlap between input images, determined by rotation step sizes ϕ , is another factor influencing reconstruction accuracy. Mesh quality is assessed by comparing the texture extraction accuracy and alignment with ground truth. Finally, a runtime analysis measures the scalability of computational costs with the number of input images and processing steps, balancing efficiency and accuracy in real-time applications. These experiments aim to understand the influence of these parameters on reconstruction quality, efficiency, and robustness.




Fig. 2. Spherical coordinates, where we refer to the polar angle θ as viewing angle, and to the azimuthal angle ϕ as

Furthermore, we assess mesh quality, focusing on the accuracy of texture extraction and the overall fidelity of the reconstructed surfaces. Finally, we conduct a runtime analysis to measure how computational cost scales with the number of input images and processing steps, balancing efficiency and accuracy in multi-view stereo reconstruction.

D. Dataset

To evaluate the capabilities of the proposed work, we utilize the MVS dataset [32] consisting of two sets of multi-view images, their camera parameters, and the ground-truth mesh models. It includes multiple views of each scene, captured from different angles to provide a diverse set of perspectives. The viewing angles are characterized by two key parameters: theta (θ), the polar angle, which defines the elevation or vertical angle from which the scene is viewed. 2.

$$\theta_{\text{cat}} \in \{30^\circ, 45^\circ, 75^\circ\} \wedge \theta_{\text{dog}} \in \{45^\circ, 90^\circ\}$$

Phi ϕ , the azimuthal angle, which represents the horizontal rotation around the scene. By capturing images across varying θ and ϕ angles, the dataset offers a comprehensive range of viewpoints. Figure 3 shows examples of the dataset and its target objects which are figurines of a cat and a dog. The images are taken with the camera by changing the height of

the camera with three different viewing angles. θ is the polar angle between the z-axis and the camera position, which we refer to as θ , as it can be seen in Figure 3.




Fig. 3. THU Multi-view stereo datasets [32] of a cat and a dog. The left side shows the input images captured from various viewpoints, while the right side displays the corresponding camera viewpoints and the target objects.

V. EVALUATION

This evaluation assesses the performance of the proposed method based on input images, focusing on the trade-off between accuracy and efficiency. Specifically, we analyze which input configurations yield the most precise reconstructions while maintaining computational efficiency.

Number of Input Images

In a first experiment, we show the impact of the number of input images on the geometric reconstruction quality. We used three different sets of images according to the three camera viewing angle θ (30° , 45° , 75°). While several combinations of rotation angles between images are possible, we chose the best result for each number of images, neglecting factors such as overlap between images as these are investigated in the following experiments. Figure 7 shows that for each of these data sets, the model converges for both CD and IoU approximately from 15 images onward while the best performance is achieved with the maximum number of input images of 36. However, some outliers deviate significantly from this trend, which can be attributed to the effects of overlap and coverage.

Runtime

In Figure 4, the model exhibits a linear runtime increase with the number of images up to 18. At 36 images, a drop-off occurs as VGGsFm is downscaled by half to fit within the available VRAM, resulting in a lower-quality reconstruction but a more accurate overall outcome. The runtime is divided into three parts: segmentation (negligible), sparse reconstruction (scales linearly with image count), and mesh extraction (consistently 4–5 minutes).

Viewing Angle Theta θ

In the first experiment, we investigate how different camera angles θ affect reconstruction quality. Therefore, the relationship between the number of images and different camera angles is used on the example of the cat and dog figurines. Figure 5 presents the best reconstruction results in terms of IoU and CD for each tested viewing angle θ . While multiple configurations are possible, we report only the optimal results for each angle.




Fig. 4. Runtime vs. input images for $\theta = 45^\circ$, $\Delta\phi = 10^\circ$

Our findings indicate that the ideal angle of incidence for both objects is 45° . At this angle, VGGsFm achieves the most effective feature extraction, as it captures both the top and front of the object within the same image. This results in a higher number of extracted points without increasing the total number of input images, outperforming alternative angles such as 30° and 75° for the cat and 90° for the dog.




Fig. 5. IoU and Chamfer Distance for cat figurine at different θ angles.

Overlap of Input Images

To reduce runtime and the computational effort required for mesh reconstruction, the following experiments explore how the input image set can be optimized. A key factor in this process is the overlap of input images. With a small rotation step ($\Delta\phi$) and few images, the visible areas are well-reconstructed, but limited coverage lowers benchmark scores despite reasonable results for small datasets. On the other hand, using few images and a large rotation step ($\Delta\phi$), leads to insufficient feature overlap and therefore an incomplete mesh. Increasing the number of images with a small $\Delta\phi$ enhances reconstruction quality, but the results are still limited to the visible areas. A balanced approach, such as 12 images with $\Delta\phi = 30^\circ$, achieves comparable quality to using three times as many images, improving both accuracy and coverage.




Fig. 6. IoU and Chamfer Distance for dog figurine at fixed θ and $\Delta\phi$

Nr. of images	$\Delta\phi$ [DEG]	CD \downarrow	IoU [%] \uparrow	Runtime [min] \downarrow
4	10°	0.0376	47.91	8.75
4	90°	0.0277	43.52	9.70
12	10°	0.0162	63.32	16.59
12	30°	0.0100	78.72	16.06
36	10°	0.0087	82.91	21.95

TABLE I

OVERLAP EXTREMES VS. IMAGE COUNT FOR CAT AT $\theta = 45^\circ$

Mesh Quality

As shown in Table II, images taken with smaller θ angles, such as top-down views, result in better texture similarities. Top-down perspectives provide a clearer view of surface details, making it easier to capture fine-grained textures. In contrast, frontal views with higher θ values lead to lower quality meshes due to occlusions and the lack of sufficient top-view coverage, making it difficult to accurately reconstruct the surface. Optimizing the pipeline and capturing images from the optimal angle leads to significantly improved SSIM and LPIPS scores, which, in human perception, translates to textures that closely resemble the ground truth.

Figurine	θ [DEG]	PSNR [dB] \uparrow	SSIM \uparrow	LPIPS \downarrow
Cat	30°	33.95	0.967	0.0373
Cat	45°	31.75	0.947	0.0546
Cat	75°	13.52	0.644	0.2516
Dog	45°	34.35	0.972	0.0353
Dog	90°	22.66	0.843	0.1195

TABLE II

TEXTURE QUALITY AT DIFFERENT θ WITH 36 IMAGES. VALUES PRINTED IN **BOLD** INDICATE THE BEST RESULT FOR EACH FIGURINE

VI. SUMMARY AND OUTLOOK

This paper analyzes how the angle of incidence and the angular distance between input images affect photogrammetric reconstruction quality. We show that an optimal balance between image count and angular separation significantly enhances mesh quality, while excessive gaps hinder feature matching. A key challenge identified is scale estimation,

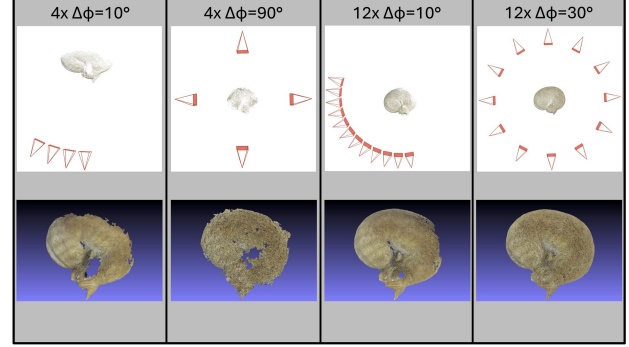


Fig. 7. Overlap of input images comparison for $\theta = 45^\circ$

which could be improved by integrating a reference object for automatic scaling. Additionally, sparse reconstruction is a major computational bottleneck, suggesting the need for more efficient alternatives. Future work should focus on optimizing reconstruction pipelines to improve runtime and scale consistency, making the approach more suitable for large-scale or real-time applications.

ACKNOWLEDGMENT

We gratefully acknowledge the support of the EU-program EC Horizon 2020 for Research and Innovation under project No. I 6114, project iChores.

REFERENCES

- [1] G. Wang, F. Manhardt, F. Tombari, and X. Ji, “Gdr-net: Geometry-guided direct regression network for monocular 6d object pose estimation,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 16 611–16 621.
- [2] Y. Su *et al.*, “Zebropose: Coarse to fine surface encoding for 6dof object pose estimation,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 6738–6748.
- [3] K. Park, T. Patten, and M. Vincze, “Pix2pose: Pixel-wise coordinate regression of objects for 6d pose estimation,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 7668–7677.
- [4] Y. Xiang, T. Schmidt, V. Narayanan, and D. Fox, “Posecnn: A convolutional neural network for 6d object pose estimation in cluttered scenes,” *arXiv preprint arXiv:1711.00199*, 2017.
- [5] D. Bauer, P. Hönig, J.-B. Weibel, J. García-Rodríguez, M. Vincze, *et al.*, “Challenges for monocular 6d object pose estimation in robotics,” *IEEE Transactions on Robotics*, 2024.
- [6] P. Ausserlechner, D. Habegger, S. Thalhammer, J.-B. Weibel, and M. Vincze, “Zs6d: Zero-shot 6d object pose estimation using vision transformers,” in *2024 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, 2024, pp. 463–469.

- [7] S. Thalhammer, J.-B. Weibel, M. Vincze, and J. Garcia-Rodriguez, “Self-supervised vision transformers for 3d pose estimation of novel objects,” *arXiv preprint arXiv:2306.00129*, 2023.
- [8] P. Hönig, S. Thalhammer, and M. Vincze, *Improving 2d-3d dense correspondences with diffusion models for 6d object pose estimation*, 2024. arXiv: 2402.06436 [cs.CV]. [Online]. Available: <https://arxiv.org/abs/2402.06436>.
- [9] Q. Shen, X. Yang, and X. Wang, “Anything-3d: Towards single-view anything reconstruction in the wild,” *arXiv preprint arXiv:2304.10261*, 2023.
- [10] G. Qian *et al.*, “Magic123: One image to high-quality 3d object generation using both 2d and 3d diffusion priors,” *arXiv preprint arXiv:2306.17843*, 2023.
- [11] A. Guédon and V. Lepetit, *Sugar: Surface-aligned gaussian splatting for efficient 3d mesh reconstruction and high-quality mesh rendering*, 2023. arXiv: 2311.12775 [cs.GR]. [Online]. Available: <https://arxiv.org/abs/2311.12775>.
- [12] X. Long *et al.*, “Wonder3d: Single image to 3d using cross-domain diffusion,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2024, pp. 9970–9980.
- [13] A. Kirillov *et al.*, “Segment anything,” *arXiv:2304.02643*, 2023.
- [14] S. Liu *et al.*, “Grounding dino: Marrying dino with grounded pre-training for open-set object detection,” *arXiv preprint arXiv:2303.05499*, 2023.
- [15] J. Wang, N. Karaev, C. Rupprecht, and D. Novotny, *Vggsfm: Visual geometry grounded deep structure from motion*, 2023. arXiv: 2312.04563 [cs.CV]. [Online]. Available: <https://arxiv.org/abs/2312.04563>.
- [16] F. Li *et al.*, “Nerf-pose: A first-reconstruct-then-regress approach for weakly-supervised 6d object pose estimation,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 2123–2133.
- [17] J. Choi *et al.*, “Tmo: Textured mesh acquisition of objects with a mobile device by using differentiable rendering,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 16 674–16 684.
- [18] P. Dai, J. Xu, W. Xie, X. Liu, H. Wang, and W. Xu, “High-quality surface reconstruction using gaussian surfels,” in *ACM SIGGRAPH 2024 Conference Papers*, Association for Computing Machinery, 2024.
- [19] B. Poole, A. Jain, J. T. Barron, and B. Mildenhall, “Dreamfusion: Text-to-3d using 2d diffusion,” *arXiv preprint arXiv:2209.14988*, 2022.
- [20] R. Liu, R. Wu, B. Van Hoorick, P. Tokmakov, S. Zakharov, and C. Vondrick, “Zero-1-to-3: Zero-shot one image to 3d object,” in *Proceedings of the IEEE/CVF international conference on computer vision*, 2023, pp. 9298–9309.
- [21] J. Xu, W. Cheng, Y. Gao, X. Wang, S. Gao, and Y. Shan, “Instantmesh: Efficient 3d mesh generation from a single image with sparse-view large reconstruction models,” *arXiv preprint arXiv:2404.07191*, 2024.
- [22] A. Radford *et al.*, *Learning transferable visual models from natural language supervision*, 2021. arXiv: 2103.00020 [cs.CV]. [Online]. Available: <https://arxiv.org/abs/2103.00020>.
- [23] R. Chen *et al.*, “Clip2scene: Towards label-efficient 3d scene understanding by clip,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 7020–7030.
- [24] Y. Zhang *et al.*, “Recognize anything: A strong image tagging model,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 1724–1732.
- [25] L. Medeiros, *Langsegmentanything: A repository for segmentation using language models*, <https://github.com/luca-medeiros/lang-segment-anything/tree/0a12766ced0503f16ee50e6fa99a9a4f5fbd4ea5>, Accessed: 2025-02-08, 2023.
- [26] J. L. Schönberger and J.-M. Frahm, “Structure-from-motion revisited,” in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [27] A. Muntoni and P. Cignoni, *PyMeshLab*, Jan. 2021. DOI: 10.5281/zenodo.4438750.
- [28] F. Williams, *Point cloud utils*, <https://www.github.com/fwilliams/point-cloud-utils>, 2022.
- [29] P.-S. Wang, Y. Liu, and X. Tong, “Dual octree graph networks for learning adaptive volumetric shape representations,” *ACM Transactions on Graphics (SIGGRAPH)*, vol. 41, no. 4, 2022.
- [30] Z. Wang, A. Bovik, H. Sheikh, and E. Simoncelli, “Image quality assessment: From error visibility to structural similarity,” *IEEE Transactions on Image Processing*, vol. 13, no. 4, pp. 600–612, 2004. DOI: 10.1109/TIP.2003.819861.
- [31] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang, “The unreasonable effectiveness of deep features as a perceptual metric,” in *CVPR*, 2018.
- [32] S. SAKAI, K. ITO, T. AOKI, T. WATANABE, and H. UNTEN, “Phase-based window matching with geometric correction for multi-view stereo,” *IEICE Transactions on Information and Systems*, vol. E98.D, no. 10, pp. 1818–1828, 2015. DOI: 10.1587/transinf.2014EDP7409.