

A Comparison of Tiny-Nerf Versus Spatial Representations for 3D Reconstruction

Saulo Abraham Gante, Juan Irving Vasquez,
Marco Antonio Valencia, Mauricio Olgún-Carbajal

Instituto Politécnico Nacional,
Centro de Innovación y Desarrollo Tecnológico en Cómputo,
Mexico

sganted1500@ipn.mx

Abstract. Neural rendering has emerged as a powerful paradigm for synthesizing images, offering many benefits over classical rendering by using neural networks to reconstruct surfaces, represent shapes, and synthesize novel views, either for objects or scenes. In this neural rendering, the environment is encoded into a neural network. We believe that these new representations can be used to codify the scene for a mobile robot. Therefore, in this work, we perform a comparison between a trending neural rendering, called tiny-NeRF, and other volume representations that are commonly used as maps in robotics, such as voxel maps, point clouds, and triangular meshes. The target is to know the advantages and disadvantages of neural representations in the robotics context. The comparison is made in terms of spatial complexity and processing time to obtain a model. Experiments show that tiny-NeRF requires three times less memory space compared to other representations. In terms of processing time, tiny-NeRF takes about six times more to compute the model.

Keywords: Neural rendering, NeRF, 3D reconstruction, mapping.

1 Introduction

The recent and continuous advances in neural rendering have shown numerous applications and became a new field of study in the graphics community. Some of these efforts are in the implicit functions which represent shapes in three dimensions (3D) [3, 15]. The tool for creating the neural representations is a multi-layer perceptron (MLP). This MLP works as a general implicit function approximator. On the other hand, there are plenty of methods to reconstruct surfaces, and represent shapes and volumes in 3D space. Some examples are meshes [2], point clouds [1], voxel maps [7], and octrees [4]. The latter offers diverse capabilities to reconstruct or create 3D models and those have diverse applications in robotics and artificial vision [18, 10, 17, 8].

In this paper, we perform a comparison between neural representations, point clouds, meshes, and voxel maps in terms of memory space and processing time required to obtain a model. The main objective is to show, clarify or put some important considerations for future works related to object reconstruction.

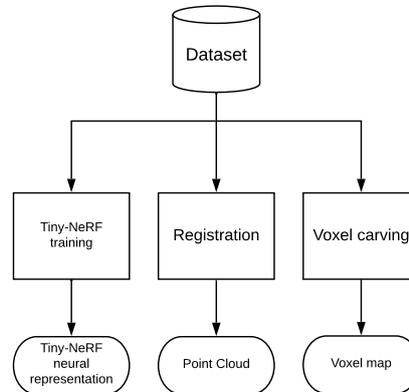


Fig. 1. General diagram of the experiments. Given a dataset which contains positions in 3D space and images, the data required is extracted with a simulator, as a result the three representations to be compared are obtained.

The MLP employed follows the architecture proposed in tiny-NeRF by the authors of Neural Radiance Fields (NeRF) [6]. We design a grid search-based experimentation. For the tiny-NeRF, the independent variables are i) learning rate ii) encoding functions and iii) seed; the resulting grid search has 36 experimental units. In addition, the experiment employs the same capturing positions to create the neural model, point cloud, and voxel map. The experiment shows that neural representations require 3 times less memory to store a model but on the other hand they takes about 6 times more to compute a model concerning the other representations.

The rest of the paper is structured as follows. Section 2 introduces the required concepts. Section 3 presents the related work and advances of volume, surface and neural representations. In Section 4, we present the methodology used to perform the experiments carry out in Section 5, where results are also reported. Finally, the conclusions and the future work are given in Section 6.

2 Preliminaries

In this section, we define certain concepts required to understand the topics tackled in this paper. A commonly used data representation is the point cloud, that is a representation of body shapes and is made by points mapped in the 3D space which are usually produced by sensors or scanners.

The point cloud could be processed in order to create more accurate representations. Meshes are representations of shapes formed by a set of nodes and connections between them, one of the advantages is that it has a range of different resolutions which means that it could be as accurate as wanted but the more resolution those have the more computation it needs to complete the representation. Another 3D shape representation is the voxel, which is a cube of unitary distance, and the union of a set creates a voxel map representation.

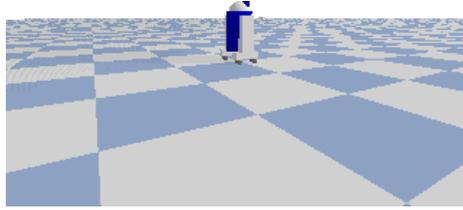


Fig. 2. World setup. An object is placed in the simulated world.

The voxel map represents shapes of objects or it is possible to represent the volume/solid by using a voxel carving technique. Those representations are commonly used to reconstruct shapes, objects and maps.

According to [16] there is no definition for neural rendering and suggests a definition for Neural Rendering as: “Deep image or video generation approaches that enable explicit or implicit control of scene properties such as illumination, camera parameters, pose, geometry, appearance, and semantic structure.”

A recent proposed neural rendering technique is NeRF [6], it became one of the most popular and extensively used to render objects in 3D space due to its capabilities to create novel views in the reconstructed scene.

NeRF [6] is an approach for creating novel view synthesis, it uses a set of input views to optimize a continuous volumetric scene function, as a result, this optimization produces a novel view of a complex scene. Its input is a 5D vector function, which contains the 3D space location (x,y,z) and 2D viewing direction (θ, ϕ) and the output is an emitted color: Red, Green, Blue (r,g,b) and volume density (σ) . NeRF uses the concept of encoding functions where the purpose of these functions is to map the input into a higher dimensional space where the MLP can more easily approximate higher frequency functions.

To generate a NeRF from a specific viewpoint, first, a set of rays are marched through the scene, the data generated is fed into the neural network and produce a set of $RGB\sigma$ values then the data is structured into a 2D image.

3 Related Work

The 3D reconstruction of spaces and objects is not a new research topic and has many approaches which reconstruct scenes employing different techniques, the accuracy of the volumetric representations relies on the resolution employed to map, and the more resolution is wanted the more computation is needed which means more time is needed to achieve a good result.

Volume representations in 3D space have many methods to represent synthetically objects, like meshes [2], point clouds [1], voxel maps [7], and octree [4]. Those offer diverse capabilities to reconstruct or create 3D models and those have diverse applications in robotics and artificial vision [8, 10, 17, 18]. Despite the popularity that they have, resolution of the representations is one of its cons. Also, the memory consumption between them is variable and usually requires memory in the order of Megabytes (MB).

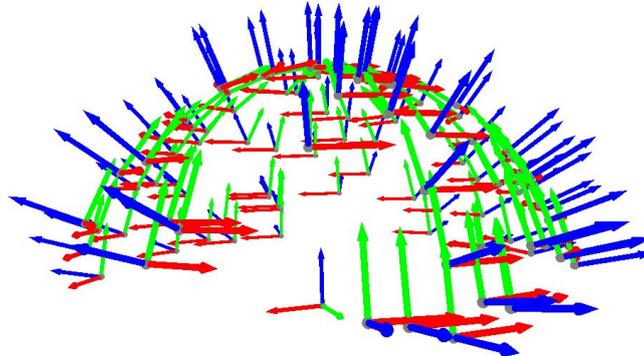


Fig. 3. Capturing poses. In this figure it is shown the 106 poses used, the frame of reference is the described by OpenGL [5] for synthetic cameras.

Neural rendering has gained popularity since it employs a multi-layer perceptron (MLP) to achieve these tasks [11, 12]. In [3, 15] are presented different techniques to represent shapes and volumes in 3D space.

The proposed methods concentrate its effort in creating those representations and compare it with state-of-the-art. On the other hand, in [14] they propose an approach for volume compression and compare it with voxel maps. A Simultaneous Localization And Mapping system is proposed by [13], they compare it with truncated signed distance function (TSDF) method, both approaches do a comparison in terms of memory consumption, and stand out a good performance.

We believe that time taken in the process of the reconstructions is an important variable to take in consideration, and that the consulted approaches do not report those differences in terms of time.

4 Methodology

We want to experiment with neural representations, exploring the advantages that those have over existent representations used for 3D space reconstruction. We do a comparison between neural representations proposed by the authors of NeRF [6], and three different spatial representations used to model objects, such as meshes, voxel maps, and point clouds.

We use a dataset that contains 106 (see Fig. 3) pairs of sensor poses, and using those poses, we extract the required data in order to create the proposed representations (Figure 1). The main reason to use one dataset is that we want to give the algorithms the same point of view for a fair comparison in terms of data that could be extracted given the poses in the data set.

4.1 Dataset

Given that we propose a comparison with synthetic data, we use a simulator to render images of an object (simulating a camera inside the simulated world). See Figure 2.

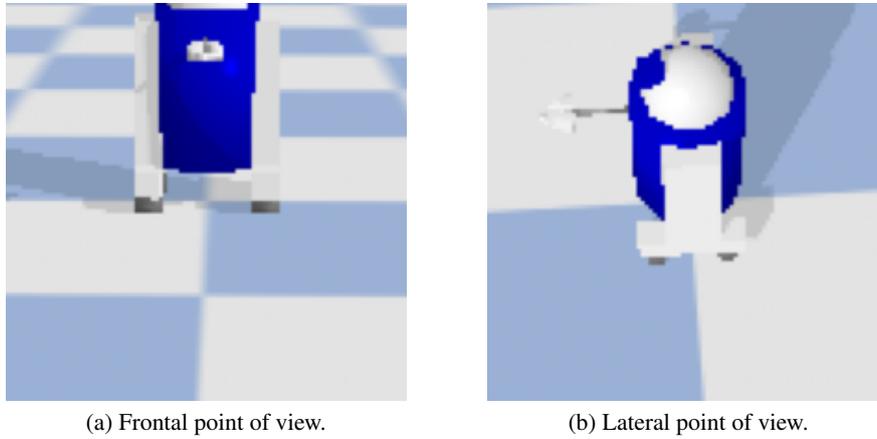


Fig. 4. Images of the object of study extracted using Pybullet.

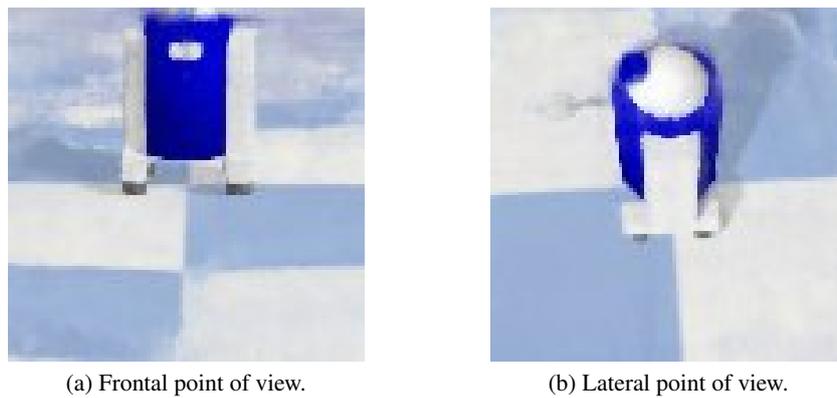


Fig. 5. Neural representations created with the Tiny-NeRF.

Then, a world is needed to set up, configure it with a ground, and place the mesh of the object of interest in it. From the synthetic datasets used in NeRF [6], we apply transformation matrices as in Equation (1):

$$T = \begin{pmatrix} R & p \\ 0 & 1 \end{pmatrix}, \quad (1)$$

where R indicates the rotation matrix, whose values represent the rotations over the three axes, and the p indicates the position vector, whose values contain the position of a body in a 3D space (x, y, z) . Please see Figure 3.

Having those positions in 3D space the required datasets are extracted, that is RGB images, ray casting points, and depth data. All in order to create the proposed reconstructions.

Table 1. Grid search. ID express an identification number, the variable values employed for each experiment with Tiny-NeRF and the results expressed in terms of Loss and PSNR.

ID	Factor 1	Factor 2	Factor 3	Metric 1	Metric 2
	Seed	Learning Rate	Coding functions	Loss	PSNR (dB)
1	2057	5×10^{-3}	6	0.5463	2.6257
2	2057	5×10^{-3}	9	0.0030	25.2288
3	2057	5×10^{-3}	10	0.0493	13.0715
4	2057	5×10^{-3}	12	0.5463	2.6257
5	2057	5×10^{-4}	6	0.5463	2.6257
6	2057	5×10^{-4}	9	0.0025	26.0206
7	2057	5×10^{-4}	10	0.5463	2.6257
8	2057	5×10^{-4}	12	0.5463	2.6257
9	2057	5×10^{-5}	6	0.5463	2.6257
10	2057	5×10^{-5}	9	0.0026	25.8503
11	2057	5×10^{-5}	10	0.5463	2.6257
12	2057	5×10^{-5}	12	0.5463	2.6257
13	7461	5×10^{-3}	6	0.0921	10.3574
14	7461	5×10^{-3}	9	0.0032	24.9485
15	7461	5×10^{-3}	10	0.5463	2.6257
16	7461	5×10^{-3}	12	0.5463	2.6257
17	7461	5×10^{-4}	6	0.5463	2.6257
18	7461	5×10^{-4}	9	0.0026	25.8503
19	7461	5×10^{-4}	10	0.5463	2.6257
20	7461	5×10^{-4}	12	0.5463	2.6257
21	7461	5×10^{-5}	6	0.5463	2.6257
22	7461	5×10^{-5}	9	0.0025	26.0206
23	7461	5×10^{-5}	10	0.5463	2.6257
24	7461	5×10^{-5}	12	0.5463	2.6257
25	5680	5×10^{-3}	6	0.5463	2.6257
26	5680	5×10^{-3}	9	0.0032	24.9485
27	5680	5×10^{-3}	10	0.5463	2.6257
28	5680	5×10^{-3}	12	0.0033	24.8149
29	5680	5×10^{-4}	6	0.5463	2.6257
30	5680	5×10^{-4}	9	0.0027	25.6864
31	5680	5×10^{-4}	10	0.5463	2.6257
32	5680	5×10^{-4}	12	0.0024	26.1979
33	5680	5×10^{-5}	6	0.5463	2.6257
34	5680	5×10^{-5}	9	0.0027	25.6864
35	5680	5×10^{-5}	10	0.5463	2.6257
36	5680	5×10^{-5}	12	0.0027	25.6864

4.2 Point Cloud and Voxel Map

Open3D library [21] allows us to visualize objects and create representations. For the point cloud, it is created by the use of ray-tracing which emits synthetic rays in simulation when those touches or intersect with a surface return a value, having this is possible to calculate in \mathbb{R}^3 and map them into points in space, creating a point cloud.

Table 2. Average of the results in grid search.

LR	Functions	Loss	PSNR (dB)
5×10^{-3}	6	0.3949	4.0351
5×10^{-3}	9	0.0031	25.0863
5×10^{-3}	10	0.3806	4.1953
5×10^{-3}	12	0.3653	4.3735
5×10^{-4}	6	0.5463	2.6256
5×10^{-4}	9	0.0026	25.8502
5×10^{-4}	10	0.5463	2.6256
5×10^{-4}	12	0.365	4.3770
5×10^{-5}	6	0.5463	2.6256
5×10^{-5}	9	0.0026	25.8502
5×10^{-5}	10	0.5463	2.6256
5×10^{-5}	12	0.3651	4.3758

This process is repeated every capture, then the resulting points are concatenated and filtered to reduce possible noise created by captures. We create a voxel model using the technique of *voxel carving*, using a pinhole camera and homogeneous transformation matrix is possible to create a voxel dense given the resulting images and a silhouette to employ a *carve silhouette* method provided by Open3D, resulting in a voxel model.

4.3 Tiny-NeRF

As explained above, NeRF [6] receives as input a set of data that express location and viewing direction where the output is an emitted color and a volume density.

Tiny-NeRF is a simplified version of NeRF, which is an MLP conformed by 6 fully-connected ReLU layers each with 256 filter size, one fully-connected ReLU layer with a filter size of 64 then an output layer that expresses the emitted $RGB\sigma$ at a certain position with a four filter size layer. The process starts by getting rays according to the pose, then the returned rays become useful to map 3D points which are going to be fed into Tiny-NeRF input, the output of the model is used to compute opacities and RGB data, finally the weights are calculated and the process is repeated.

5 Experiments

We evaluate the Tiny-NeRF describing a grid search where certain variables are changed over the experiments. Using the same position captures we perform reconstruction with voxels and a point cloud. All the data was synthetic and obtained using Open3D.

Our experiments run in Python and the libraries employed are Pytorch [9] for the MLP or neural representations, Pybullet and Open3D [21]. The hardware employed for those experiments is the CPU/GPU provided by Colab which allows us to use a Graphic card: Tesla P100-PCIE-16GB with 16GB of GPU-RAM, 25.46 GB of RAM, and 166.83 GB in Hard disc drive.

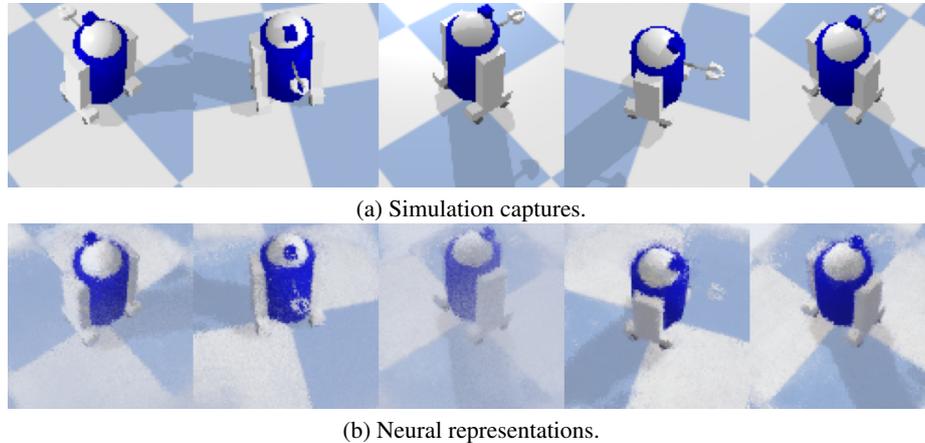


Fig. 6. Qualitative results using Tiny-NeRF. We extracted different poses and visualizations.

5.1 Tiny-NeRF Training

Tiny-NeRF is a simplified version of NeRF, which is an MLP conformed by six fully-connected ReLU layers each with a 256 filter size, one fully-connected layer with a filter size of 64 then an output that expresses the $RGB\sigma$ values. The grid search proposed to vary over three variables and the values are:

- **Seed:** 2057, 5680 and 7461,
- **Learning rate:** $5x10^{-3}$, $5x10^{-3}$ and $5x10^{-3}$,
- **Encoding functions :** 6, 9, 10 and 12.

For the Neural Networks (NN) training a commonly used metric is Loss since it evaluates how bad predicts on an example, the *Peak Signal-to-Noise Ratio* (PSNR) is used to measure the ratio between a signal and the noise which affects the representation of this signal; in this case, the PSNR is used to measure how well the Tiny-NeRF does a representation compared to the original images. On the other hand, to measure time the unit employed is seconds (s) and to measure space in memory we utilize MB.

To obtain the data set we employed Pybullet simulator which let us set simulated worlds, set objects in it (Figure 2) and create pinhole cameras to extract or create synthetic images, among other things. Once the object is set in the world, it is possible to create a synthetic camera given its position, target position, field of view (FOV), near and far plane distance, weight and height of the image. The positions are given by the data capturing positions, the FOV is 17.70° , the weight and height are equal to 100. Resulting in images like the ones in Figure 4.

The experiments with the Tiny-NeRF are iterated for five thousand epochs each and the variables are modified at each experiment, the number of experiments is 36. The results of this experiment are shown in Table 1. The entire experiment took about 65,100 seconds which means that approximately every experiment took 1,808.33 seconds to be completed. To summarize the information in Table 1, the average was calculated (Table 2) in order to easily extract which parameters perform better results with the MLP.

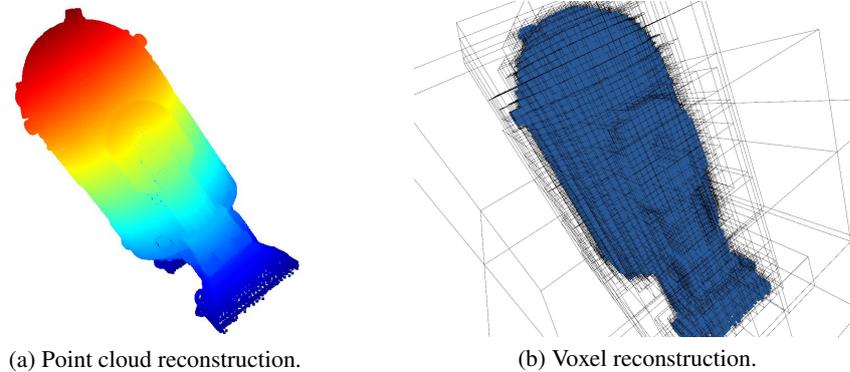


Fig. 7. Volumetric reconstructions.

Analysis of experiments showed that nine coding functions help the Tiny-NeRF to accurately (Figure 5) create a neural representation of the object, and the learning rate helped to achieve good performance in fewer epochs. Additionally, the neural representations took **1,808.33 seconds** to complete an experiment, and the memory space to store a representation is **1.5 MB**.

Additionally to PSNR, we perform evaluations over two more metrics SSIM and LPIPS [19, 20] which are commonly used to measure distances over images, looking for a measure of how well the Tiny-NeRF is rendering views. Comparing images like the ones in Fig. 6 the metrics proposed gave as a result **0.8481** and **0.0565**, respectively. Those results affirm that the representations are good in quality but it could improved.

5.2 Comparison of Tiny-NeRF Versus Spatial Representations

Once Tiny-NeRF has been trained and the representations were created, we compared the time taken to do a representation. To measure the time, it was printed every time a process started and finished the difference between those shows the time taken. The space in memory is measured by the file space in memory that is required to store the representations.

The data employed to reconstruct was obtained by capturing in the positions of the data set, mentioned above, once the captures are done the process of data was done employing Open3D [21].

The point cloud (Figure 7 (a)) was obtained by mapping the points resultant of a ray-tracing operation into XYZ or 3D space, those points are concatenated and finally filtrated to avoid noise in the reconstruction. The experiment took about **2 seconds** and the memory space needed is **12 MB**.

The resultant voxel map (Figure 7 (b)) was created with the voxel carving method which not only reconstructs the surface of an object, it creates a voxel map that is a cube of certain dimensions and according to the visualized data, the algorithm carves the shape into it, creating a solid voxel representation. The experiment took about **166 seconds** and the memory space needed is **21.2 MB**.

The performed experiments results showed that implicit or neural representation requires at least 3 times less memory compared with other representations (Table 3).

Table 3. Comparative of memory size.

Representation	Size (MB)
Meshes	4.5
Point cloud	12.0
Voxelization	21.2
Implicit representation	1.5

Table 4. Comparative of time taken to perform a representation.

Representation	Time (s)
Meshes	28800
Point cloud	2
Voxelization	166
Implicit representation	1008

In terms of time to process a representation, point clouds and voxel maps build the representations in about 6 times less time than the implicit representations (Table 4).

6 Conclusions and Future Work

This paper tackles the trend research topic, neural rendering, which has many advances in graphics generation. We compare a simplified version of NeRF with different volume representations commonly used in robotics and vision reconstruction tasks, all compared in terms of memory space and time to build representations.

First, we experimented with Tiny-NeRF that computes the colors over a certain position with a viewing direction; the experiments were conducted by a grid search looking to perform good representations of an object. In addition, we perform a reconstruction using voxels and point clouds. The comparison, in terms of memory space and time, shows that the Tiny-NeRF architecture (MLP) requires less memory but takes more time to build a representation.

On the other hand, this experiment showed that the neural representation relies on the fine-tuning of the variables implied in the training of the MLP. We believe that the results of the experiments can offer some relevant information or considerations to take when a reconstruction task is needed. In a future work we will experiment with more objects and with a mobile manipulator robot.

References

1. Berger, M., Tagliasacchi, A., Seversky, L., Alliez, P., Levine, J., Sharf, A., Silva, C.: State of the Art in Surface Reconstruction from Point Clouds. In: Eurographics - State of the Art Reports, pp. 161–185 (2014). DOI: 10.2312/egst.20141040.
2. Delingette, H.: General Object Reconstruction Based on Simplex Meshes. International Journal of Computer Vision, vol. 32, pp. 111–146 (1999). DOI: 10.1023/A:1008157432188.
3. Genova, K., Cole, F., Sud, A., Sarna, A., Funkhouser, T.: Local Deep Implicit Functions for 3D Shape. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 4857–4866 (2020). DOI: 10.1109/CVPR42600.2020.00491.

4. Hornung, A., Wurm, K.M., Bennewitz, M., Stachniss, C., Burgard, W.: OctoMap: An Efficient Probabilistic 3D Mapping Framework Based on Octrees. *Autonomous Robots*, vol. 34, pp. 189–206 (2013). DOI: 10.1007/s10514-012-9321-0.
5. Learn OpenGL: Camera <https://learnopengl.com/Getting-started/Camera> (2018)
6. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis. In: *European Conference on Computer Vision*, vol. 12346, pp. 405–421 (2020). DOI: 10.1007/978-3-030-58452-8_24.
7. Muglikar, M., Zhang, Z., Scaramuzza, D.: Voxel Map for Visual SLAM. In: *IEEE International Conference on Robotics and Automation (ICRA)*, pp. 4181–4187 (2020). DOI: 10.1109/ICRA40945.2020.9197357.
8. Mur-Artal, R., Tardós, J.D.: ORB-SLAM2: An Open-Source SLAM System for Monocular, Stereo, and RGB-D Cameras. *IEEE Transactions on Robotics*, vol. 33, no. 5, pp. 1255–1262 (2017). DOI: 10.1109/TRO.2017.2705103.
9. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., De-Vito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., et al.: PyTorch: An Imperative Style, High-Performance Deep Learning Library. In: *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, pp. 8026–8037 (2019). DOI: 10.48550/arXiv.1912.01703.
10. Respass, V.M., Devitt, D., Fedorenko, R., Klimchik, A.: Fast Sampling-Based Next-Best-View Exploration Algorithm for a MAV. In: *IEEE International Conference on Robotics and Automation (ICRA)*, pp. 89–95 (2021). DOI: 10.1109/ICRA48506.2021.9562107.
11. Saito, S., Huang, Z., Natsume, R., Morishima, S., Kanazawa, A., Li, H.: PIFu: Pixel-Aligned Implicit Function for High-Resolution Clothed Human Digitization, pp. 2304–2314 (2019). DOI: 10.48550/arXiv.1905.05172.
12. Sitzmann, V., Martel, J.N.P., Bergman, A.W., Lindell, D.B., Wetzstein, G.: Implicit Neural Representations with Periodic Activation Functions. In: *34th Conference on Neural Information Processing System*, pp. 1–35 (2020). DOI: 10.48550/arXiv.2006.09661.
13. Sucar, E., Liu, S., Ortiz, J., Davison, A.: iMAP: Implicit Mapping and Positioning in Real-Time. In: *Proceedings of the International Conference on Computer Vision (ICCV)*, pp. 6229–6238 (2021). DOI: 10.48550/arXiv.2103.12352.
14. Tang, D., Singh, S., Chou, P.A., Häne, C., Dou, M., Fanello, S.R., Taylor, J., Davidson, P.L., Guleryuz, O.G., Zhang, Y., Izadi, S., Tagliasacchi, A., Bouaziz, S., Keskin, C.: Deep Implicit Volume Compression. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1293–1303 (2020). DOI: 10.1109/CVPR42600.2020.00137.
15. Tang, J., Lei, J., Xu, D., Ma, F., Jia, K., Zhang, L.: Sign-Agnostic CONet: Learning Implicit Surface Reconstructions by Sign-Agnostic Optimization of Convolutional Occupancy Networks. In: *International Conference on Computer Vision*, pp. 1–16 (2021)
16. Tewari, A., Fried, O., Thies, J., Sitzmann, V., Lombardi, S., Sunkavalli, K., Martin-Brualla, R., Simon, T., Saragih, J.M., Nießner, M., Pandey, R., Fanello, S.R., Wetzstein, G., Zhu, J.Y., Theobalt, C., Agrawala, M., Shechtman, E., Goldman, D.B., Zollhöfer, M.: State of the Art on Neural Rendering. *Computer Graphics Forum*, vol. 39, no. 2, pp. 701–727 (2020). DOI: 10.1111/cgf.14022.
17. Uyanik, C., Secil, S., Ozkan, M., Dutagaci, H., Turgut, K., Parlaktuna, O.: SPGS: A New Method for Autonomous 3D Reconstruction of Unknown Objects by an Industrial Robot. In: *Annual Conference Towards Autonomous Robotic Systems*, vol. 10965, pp. 15–27 (2018). DOI: 10.1007/978-3-319-96728-8_2.
18. Vasquez-Gomez, J.I., Sucar, L.E., Murrieta-Cid, R.: View Planning for 3D Object Reconstruction with a Mobile Manipulator Robot. In: *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 4227–4233 (2014). DOI: 10.1109/IROS.2014.6943158.

Saulo Abraham Gante, Juan Irving Vasquez, et al.

19. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: 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.
20. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The Unreasonable Effectiveness of Deep Features as a Perceptual Metric. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 586–595 (2018). DOI: 10.1109/CVPR.2018.00068.
21. Zhou, Q.Y., Park, J., Koltun, V.: Open3D: A Modern Library for 3D Data Processing. *Computer Vision and Pattern Recognition*, pp. 1–6 (2018). DOI: 10.48550/arXiv.1801.09847.