

HoloGS: Instant Depth-based 3D Gaussian Splatting with Microsoft HoloLens 2

Miriam Jäger*, Theodor Kapler, Michael Feßenbecker, Felix Birkelbach, Markus Hillemann, Boris Jutzi

Institute of Photogrammetry and Remote Sensing (IPF), Karlsruhe Institute of Technology (KIT), Karlsruhe, Germany
(miriam.jaeger, markus.hillemann, boris.jutzi)@kit.edu

Keywords: 3D Gaussian Splatting, Microsoft HoloLens 2, Depth Sensor, Point Cloud, 3D Reconstruction, Neural Radiance Fields

Abstract

In the fields of photogrammetry, computer vision and computer graphics, the task of neural 3D scene reconstruction has led to the exploration of various techniques. Among these, 3D Gaussian Splatting stands out for its explicit representation of scenes using 3D Gaussians, making it appealing for tasks like 3D point cloud extraction and surface reconstruction. Motivated by its potential, we address the domain of 3D scene reconstruction, aiming to leverage the capabilities of the Microsoft HoloLens 2 for instant 3D Gaussian Splatting. We present HoloGS, a novel workflow utilizing HoloLens sensor data, which bypasses the need for pre-processing steps like Structure from Motion by instantly accessing the required input data i.e. the images, camera poses and the point cloud from depth sensing. We provide comprehensive investigations, including the training process and the rendering quality, assessed through the Peak Signal-to-Noise Ratio, and the geometric 3D accuracy of the densified point cloud from Gaussian centers, measured by Chamfer Distance. We evaluate our approach on two self-captured scenes: An outdoor scene of a cultural heritage statue and an indoor scene of a fine-structured plant. Our results show that the HoloLens data, including RGB images, corresponding camera poses, and depth sensing based point clouds to initialize the Gaussians, are suitable as input for 3D Gaussian Splatting.

1. Introduction

3D scene reconstruction is a fundamental task in the fields of computer vision, computer graphics and photogrammetry. Recently, however, methods have been gaining popularity that have the potential to revolutionize the classical workflows. This has been particularly initiated by the pioneering research on Neural Radiance Fields (NeRFs) (Mildenhall et al., 2020).

NeRFs and HoloLens. NeRFs enable the rendering of novel views with the so-called view synthesis of a 3D scene in space with a neural network. The neural network is trained based on a set of images and corresponding camera poses and estimates a position-dependent density value and view-dependent RGB color values per position. Through the volume density of points in a continuous space, geometries can be extracted. However, this requires techniques such as density thresholding or other methods of generating explicit 3D (surface) reconstructions from continuous neural network outputs (Oechsle et al., 2021; Wang et al., 2021; Yariv et al., 2021; Darmon et al., 2022; Park et al., 2019; Zhang et al., 2022; Li et al., 2023; Jäger and Jutzi, 2023), while the density carries an inherent uncertainty (Jäger et al., 2023). Most commonly, traditional methods like Structure from Motion (SfM) are used to calculate the interior orientation and the camera poses needed for training the NeRFs in a pre-processing step. As an alternative to SfM, the Microsoft HoloLens has proven to be an interesting interface, since it enables the extraction of the required input data, the images and corresponding poses (Jäger et al., 2023). Moreover, it has consistently demonstrated its efficacy as a mapping system (Weinmann et al., 2020; Weinmann et al., 2021; Hou et al., 2024) and enables the real-time (Haitz et al., 2023), highly detailed, colorized, 3D scene reconstruction and mobile mapping (Jäger et al., 2023) with NeRFs.

Gaussian Splatting and HoloLens. With regard to 3D scene reconstruction, particularly 3D Gaussian Splatting (GS) (Kerbl et al., 2023) is outstanding due to its explicit representation of the scene utilizing 3D Gaussians. During optimization, these Gaussians are densified and adapted, undergoing growth, shrinkage, and adjustments in color and shape, until the photometric error between rendered images and training images becomes minimal. In contrast to the continuous radiance field representation of NeRF, these Gaussians explicitly represent the scene geometry, enabling direct access to it. When it comes to photogrammetry and 3D computer vision, this is particularly of interest for 3D point cloud extraction and surface reconstruction. In contrast to most NeRF methods, further input data in the form of a point cloud is required for 3D Gaussian Splatting, for which the sparse point cloud from SfM is usually used. This point cloud is used to initialize the Gaussians. Thus, pre-processing and calculation steps are required to compute the camera poses and sparse point cloud from the images. At this point, the Microsoft HoloLens once again becomes relevant. Alongside the RGB images and their corresponding camera poses, the HoloLens provides depth images and corresponding camera poses, which can be transformed into the required point cloud. This enables instant 3D scene reconstruction with 3D Gaussian Splatting, i.e. without additional time-consuming pre-processing steps.

In this work, we present HoloGS (Figure 1), for an instant 3D scene reconstruction by 3D Gaussian Splatting with Microsoft HoloLens 2 data. This is done directly based on sensor information, since the HoloLens enables to access the required input data, i.e. the images, camera poses and the point cloud from depth sensing in real-time. We investigate whether the data quality of the HoloLens is sufficient for 3D Gaussian Splatting. In order to evaluate our workflow, we additionally follow the traditional pipeline, which uses COLMAP (Schönberger and Frahm, 2016), to estimate the camera poses and the sparse point cloud. Our analysis includes examining the training process

* Corresponding author

and evaluating the rendering quality of our results by using the Peak Signal-to-Noise Ratio (PSNR) and photometric loss. Furthermore, we report the geometric 3D accuracy quantitatively as well as qualitatively of the resulting densified point clouds from the Gaussian centers using cloud-to-cloud Chamfer Distance. Thereby, we envision a refilling of the (sparse) input point cloud, comparable to a post-processing step by Multi-View Stereo.

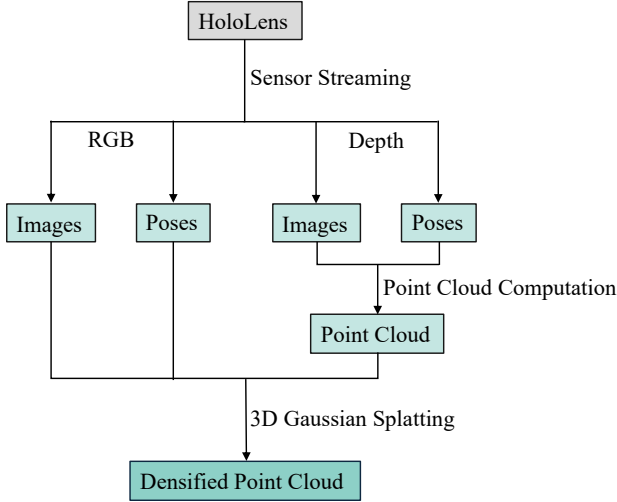


Figure 1. Flowchart of HoloGS: Via the HoloLens 2 sensor streaming, the required data is directly extracted and processed during data capturing. From the depth data, a point cloud is instant created which, together with the RGB images and corresponding camera poses, is then fed into 3D Gaussian Splatting.

We demonstrate that HoloGS with Microsoft HoloLens 2 data, comprising the RGB images with corresponding camera poses, and the point cloud from depth sensing, is suitable as input for 3D Gaussian Splatting. The rendered images reasonably reflect the geometry and appearance. Furthermore, HoloGS enables refilling by the extraction of a densified point cloud from Gaussian centers.

2. Methodology

Section 2.1 outlines the principles of the methods used to determine the input data for 3D Gaussian Splatting: via the standard method with external data from SfM and via our approach for instant 3D Gaussian Splatting with internal data from Microsoft HoloLens 2. Subsequently, Section 2.2 presents the implementation details for Gaussian Splatting. Finally, Section 2.3 describes our method for extracting the densified point cloud from Gaussian Splatting after the training process.

2.1 Initialization

External SfM data. As mentioned, the standard workflow uses SfM to determine the camera poses and the sparse point cloud in a pre-processing step. SfM in general describes the procedure of the reconstruction of a 3D scene from a set of images, which taken from different directions and positions by a camera in motion. It relies on the calculation and matching of point correspondences within an image sequence from overlapping images. Most commonly by using methods such as SIFT (Lowe, 2004). The resulting products are the camera

poses, the camera intrinsics as well as a sparse point cloud from the point correspondences. In addition, the fundamental SfM sparse point cloud allows a following Multi-View Stereo (MVS) pipeline (Schönberger and Frahm, 2016) regarding a dense reconstruction, which densifies the SfM point cloud. We consider the MVS dense reconstruction as a reference point cloud for comparison. In this paper, the (incremental) SfM technique by (Schönberger and Frahm, 2016) from the original implementation of 3D Gaussian Splatting¹ is used for the external data calculation of the camera poses and the sparse point cloud. This process is conducted using the same HoloLens RGB images to ensure uniform conditions similar to those of the following internal data approach.

Internal HoloLens data. For an instant 3D scene reconstruction with 3D Gaussian Splatting directly from Microsoft HoloLens 2, HoloGS targets three main steps, according to Figure 1: Sensor streaming, real-time point cloud computation, and instant 3D Gaussian Splatting. The HoloLens 2 server application (Dibene and Dunn, 2022) is used for requesting the data in the Microsoft HoloLens 2². The system provides access to all of the HoloLens 2 sensors, including the RGB images from the 1920×1080 camera, interior orientation (camera intrinsics) and corresponding camera poses, as well as the depth sensor for depth images and corresponding poses. Firstly, through the sensor streaming of the HoloLens, the RGB images and corresponding camera poses including the interior orientation as well as the depth images and their camera poses are queried and extracted from the sensor system during data acquisition. Secondly, the depth images are each transformed into a 3D point cloud by calculating the corresponding 3D point for each pixel in the depth image based on the interior orientation of the depth camera and the depth information. These point clouds are subsequently merged into a joint point cloud via the camera poses. Lastly, the required data, i.e., the RGB images with corresponding camera poses and the point cloud from the depth information transformed to the coordinate system of the RGB camera, is fed to 3D Gaussian Splatting as initial data. In this context, the RGB images serve as training data for optimizing the Gaussians by minimizing the photometric error between rendered images and their real counterparts at the same camera poses. The 3D points of the point cloud from the depth information forms the centers of the initial Gaussians. The other parameters of the Gaussians are initialized as in the original implementation of Kerbl et al., i.e. isotropic Gaussians with axes equal to the mean of the distance to the closest three points.

2.2 3D Gaussian Splatting Implementation

After initialization, 3D Gaussian Splatting is processed according to the original implementation. We train on the default parameters with learning rates of 0.0025 for spherical harmonics features, 0.05 for opacity adjustments, 0.005 for scaling operations and 0.001 for rotation transformations, while the training incorporates 30 000 iterations on a NVIDIA RTX3090 GPU. The photometric loss for the optimization is given by the following loss function 1.

$$L = (1 - \lambda)L_1 + \lambda L_{D-SSIM} \quad (1)$$

with $\lambda = 0.2$ by default, L_1 -Norm of the per pixel color difference and L_{D-SSIM} -Term (Kerbl et al., 2023).

¹ <https://github.com/graphdeco-inria/gaussian-splatting> (last access 12/01/2023)

² <https://www.microsoft.com/en-us/hololens/hardware> (last access 02/08/2024)

2.3 3D Point Cloud Extraction

As mentioned above, the point clouds from the external SfM or from the internal depth information of the HoloLens serve as initial Gaussians. Through the optimization process of Gaussian Splatting, additional points are generated based on color information of the images. In doing so, Gaussians grow, split, shrink or are removed. Based on this, we envision a refilling of the sparse input point cloud, comparable to a post-processing step by MVS. Especially, since the SfM sparse point cloud in particular is only created based on the features from SIFT point correspondences, it is, as the name implies, relatively sparse. After the training of 3D Gaussian Splatting, the densified and optimized point cloud can be extracted. On the assumption that 3D information containing color also exists as actual geometry in the scene, we consider the centers of the Gaussians, which represent each the mean of each Gaussian, as 3D geometry in our approach.

3. Dataset

Our experiments are based on two datasets that we captured with Microsoft HoloLens 2: An outdoor scene of a cultural heritage statue, called 'Denker' (Figure 2), and an indoor scene, called 'Ficus' of a fine-structured plant. Both datasets contain 111 images each of size 1920×1080 with camera poses, captured in a hemispherical camera framing, as well as the respective depth maps and poses, whereby we mask out depth information above 3m. The respective initialization type (Section 2.1) of the input data results in different point clouds: The SfM sparse point cloud (Figure 3(a)) including over 40 000 points, and the point cloud (Figure 3(b)) from the depth images from HoloLens including over 620 000 points .

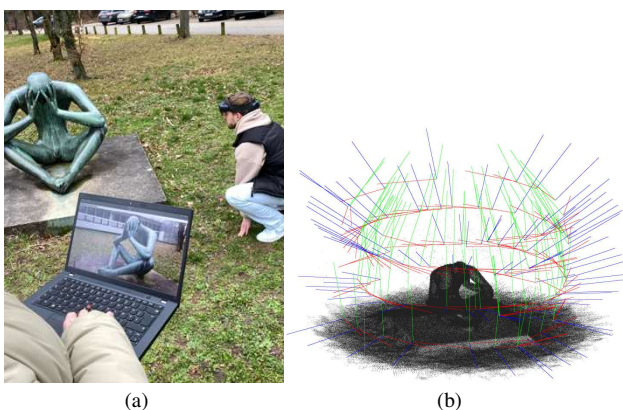


Figure 2. (a) Data capturing with Microsoft HoloLens 2 and its streaming application (Dibene and Dunn, 2022). (b) Point cloud based on depth data of the scene 'Denker' and camera poses visualized by colored coordinate frames.

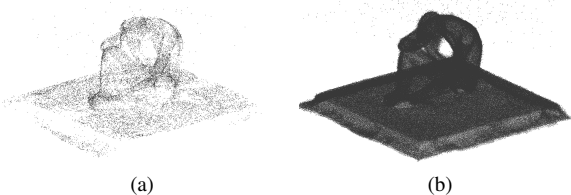


Figure 3. Initialization input 3D point clouds. (a) sparse point cloud from SfM and (b) point cloud calculated based on the depth images from HoloLens.

4. Experiments and Results

In this section, we present our experiments and results by a quantitative evaluation on analyzing the training process by rendering quality in Section 4.1. This is followed by a qualitative analysis of the rendered images in Section 4.2. Finally, we evaluate the geometric 3D reconstruction based on the densified point clouds from the Gaussian centers quantitative as well as qualitative in Section 4.3.

4.1 Training

We evaluate the training process with the Peak Signal-to-Noise Ratio (PSNR) (Mildenhall et al., 2020), which is a common metric in NeRF context. Figure 4 shows the change in PSNR and training loss (Kerbl et al., 2023) (Equation 1) over the iterations. It demonstrates that HoloGS with internal HoloLens data, including RGB images, corresponding camera poses, and point clouds derived from the depth map, leads to relatively smooth convergence of 3D Gaussian Splatting. Convergence occurs after approximately 25,000 iterations, reaching a maximum PSNR (Table 1) of around 20.55 dB for the scene 'Denker' and 20.17 dB for the scene 'Ficus'. Notably, the convergence is rapidly achieved with the HoloLens data for both scenes. In contrast, utilizing external SfM data yields higher PSNR values of 27.54 dB for the scene 'Denker' and 26.21 dB for the scene 'Ficus'. Conversely, the loss for the internal HoloLens data during training is higher than for the external SfM data for both scenes. Interestingly, the curves show peaks every 3000 iterations up to iteration 15000. These peaks can be explained by the density moderation technique of Gaussian Splatting. This technique sets the opacity values close to zero every 3000 iterations to prevent the method from getting stuck with floaters close to the camera poses which could cause an unjustified increase in the density of Gaussians (Kerbl et al., 2023).

	External SfM data	Internal HoloLens data
Denker	27.54	20.17
Ficus	26.21	20.55

Table 1. Peak Signal-to-Noise Ratio (PSNR) \uparrow in dB after 30 000 iterations for the scenes 'Denker' and 'Ficus' each with external SfM and internal HoloLens data.

4.2 Rendering Quality

The results of the rendered images in Figure 5 closely correspond to the numerical results obtained during the training process. Particularly, the external SfM data computed during pre-processing demonstrates significantly improved performance. For the scene 'Denker', HoloGS produces satisfactory results for the statue itself, comparable to those derived from SfM data. Furthermore, in scene 'Ficus', HoloGS struggles to accurately represent the fine structures of the plant, leading to noise. Overall, HoloGS generally results in blurry edges of objects in the scene. In addition, large, foggy and blurry floater artifacts can be recognized in these scenes. In contrast, for the external SfM data, these artifacts are only evident in a limited number of areas, e.g. above the head of the statue in the scene 'Denker' and in the unobserved area of the scene 'Ficus' on the ceiling.

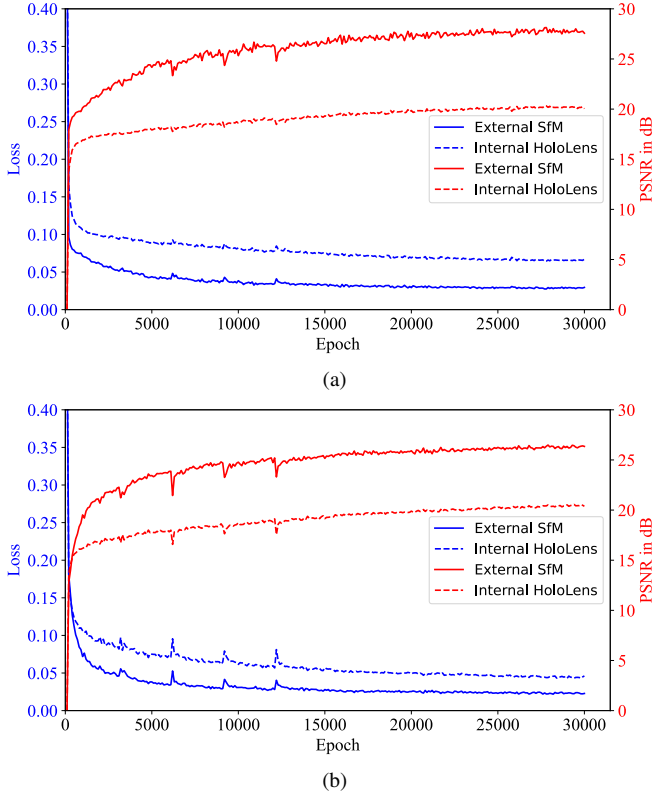


Figure 4. Comparison of the Peak Signal-to-Noise Ratio (PSNR) \uparrow in dB and loss \downarrow during the training processes with 30 000 iterations with 3D Gaussian Splatting with different types of input data. Top: (a) external SfM and internal HoloLens data on scene 'Denker'. Bottom: (b) external SfM and internal HoloLens data on scene 'Ficus'. The red curves show the PSNR, the blue curves the training loss.

4.3 3D Point Cloud Extraction

We extract the densified point cloud, which is generated during the training process like described in Section 2.3. The extracted densified point clouds of the Gaussians (Figure 6) illustrate the differences in Gaussian Splatting for 3D scene reconstruction between the external SfM data and the internal HoloLens data. For the scene 'Denker', there is an overall good coverage with the SfM data, where the structure of the object is clearly visible with sharp edges. Nonetheless, some gaps in the point cloud are evident, particularly on the platform and on the statue's arms and legs. In these areas the color differs from the rest and appears more homogeneous and low-textured. When using the internal HoloLens data, there is an identifiable structure of the object through the centers of the Gaussians. Although, the point cloud is overall noisy, with indistinct and fuzzy object edges. Additionally, large artifacts are present in the point clouds, corresponding to the floater artifacts in the rendered images. For the scene 'Ficus', a similar pattern emerges. The point cloud from external SfM data exhibits clear and sharp edges, accurately capturing the fine structure of the vegetation of the plant. However, gaps are noticeable, especially in the pot area, characterized by a uniform, low-textured color. Additionally, small floater artifacts appear intermittently above the object, particularly in areas with lower scene coverage from the captures. When using the HoloLens data, the structure of the object is also clearly recognizable, with a high coverage. Nonetheless, the point cloud remains heavily noisy, with numerous floater

artifacts, especially in areas above the object. Again, the low-textured pot of the plant is not full reconstructed.

The visual appearance of the densified point clouds is further evaluated quantitatively and qualitatively by their geometric 3D accuracy (Jensen et al., 2014) using the Chamfer cloud-to-cloud Distance from the point cloud to the reference from MVS. Consistent with the training and rendering results, the extracted point cloud exhibits similar geometric characteristics, as shown quantitatively in Table 2 and qualitatively in Figures 7 and 8. The quantitative results in Table 2 highlight clear differences in the geometric accuracy of the extracted densified point clouds, especially regarding the two types of initial input data. For the scene 'Denker', with a mean accuracy of 0.021 and a standard deviation of 0.061 for external SfM data, contrasting with a significantly lower accuracy of 0.298 and a standard deviation of 0.534 for internal HoloLens data. Similar results are obtained for the scene 'Ficus', where the use of external SfM data shows a mean geometric accuracy of 0.045, with a standard deviation of 0.261. In contrast, the use of internal HoloLens data again results in a significantly lower accuracy of 0.596, with a standard deviation of 0.891. The point clouds, illustrating the Chamfer Distances, visually underscore the quantitative findings, as shown in Figure 7 for the scene 'Denker'. It is evident that the external SfM data yields a high accuracy for the statue's surface, although a lower geometric accuracy is noticeable on low-textured, homogeneous areas. The same trend is seen for the internal HoloLens data, where lower geometric accuracy is reflected in noisy edges and floater artifacts. Figure 8 shows a similar pattern for the scene 'Ficus'. The SfM data performs well in capturing details, while the HoloLens data shows reduced accuracy, especially on fine-structured object parts like the branches.

	External SfM data		Internal HoloLens data	
	mean	std	mean	std
Denker	0.021	0.061	0.298	0.534
Ficus	0.045	0.261	0.596	0.891

Table 2. Geometric 3D accuracy via Chamfer Distance \downarrow . Mean distance (mean) and standard derivation (std) for the scenes 'Denker' and 'Ficus' each with external SfM and internal HoloLens data. Note that the reference point clouds and therefore the Chamfer Distances are non-metrical.

5. Discussion

In this paper, we introduce HoloGS to investigate the application of instant 3D Gaussian Splatting using data from the internal sensors of the Microsoft HoloLens 2. Specifically, the input data consists of RGB images with corresponding camera poses and a point cloud from the depth data of the HoloLens as initial Gaussian centers. We have shown both quantitatively and qualitatively that the internal HoloLens data is suitable for this application as it enables convergence of the Gaussian Splatting optimization process. The optimization of the training process based on the rendered RGB images converges quickly and reaches a maximum PSNR value of 20.17 for the scene 'Denker' and 20.55 for 'Ficus'. This convergence enables the rendering of novel synthetic images from different views which represent the scene visually well. Additionally, the optimization of the Gaussians during the training process enables the extraction of the densified point cloud from the Gaussian centers.

Nevertheless, limitations exist, as the results of the externally preprocessed SfM data outperform the results of the internal

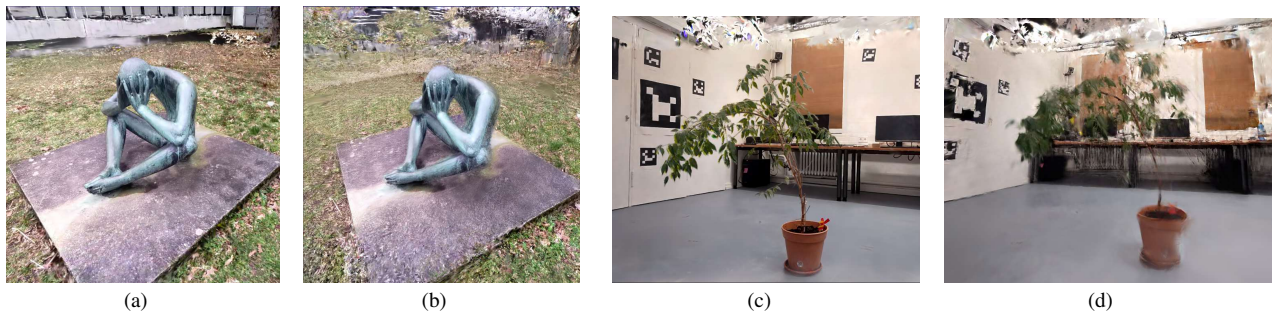


Figure 5. Rendered images. From left to right: (a) external SfM data and (b) internal HoloLens data on scene 'Denker', as well as (c) external SfM data and (d) internal HoloLens data on scene 'Ficus'.

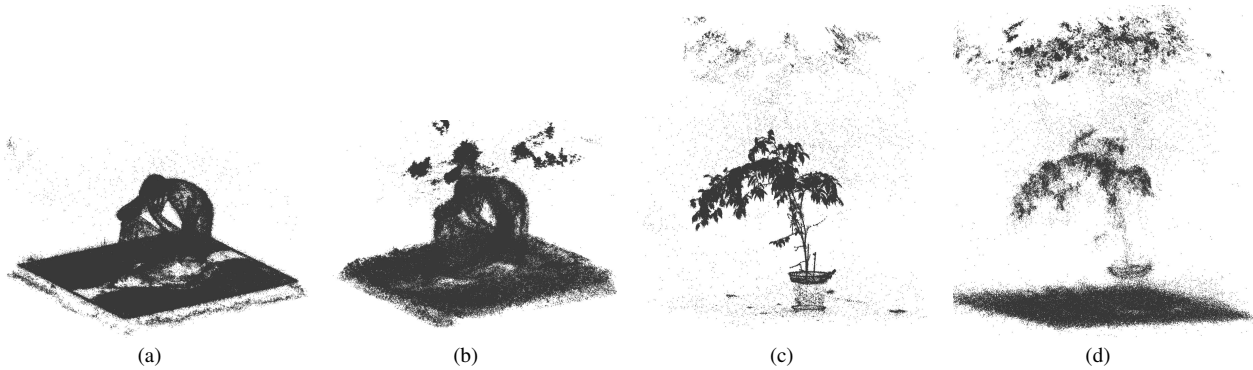


Figure 6. Point clouds by extracting the center of each Gaussian after training process. From left to right: (a) external SfM data on scene 'Denker', (b) internal HoloLens data on scene 'Denker', (c) external SfM data on scene 'Ficus', (d) internal HoloLens data on scene 'Ficus'.

HoloLens data. Thereby, a higher maximum PSNR is reached during the training, and the rendered images appear less blurry and contain less floater artifacts in comparison to the internal HoloLens data. In addition, the geometric accuracy of the 3D reconstruction of the densified point clouds with HoloLens performs 10 times weaker on average, although this may be due to floater artifacts, which weigh heavily. We suspect the cause of these discrepancies lies in the less precise camera poses of the RGB images of the internal HoloLens data, leading to blurry results and artifacts in the rendered images as well as the densified point clouds. Moreover, it could be assumed that the initial point cloud from the depth sensor may not match the correct positions of the RGB images in the coordinate system. However, since the PSNR does not continue to increase during training as Gaussians grow, shrink or are removed, we reject this assumption as a potential cause.

Therefore, firstly, considering the 3D mapping aspect, the usage of HoloLens solely as a 3D mapping system, without the 3D scene reconstruction aspect of computer graphics, a direct usage of the depth maps and the resulting point cloud from the HoloLens sensor data can be considered. This results in a point cloud with high point density and fine details, as shown by the input data (Section 3) of the internal HoloLens data approach. Secondly, from the aspect of computer graphics, we nonetheless consider the combination of HoloLens and Gaussian Splatting to be suitable. If, as suspected, the weaker results are due to the RGB camera poses, just as with the HoloLens-NeRF combination (Jäger et al., 2023), we propose the optimization of RGB camera poses during the training process. A strategy previously employed in the context of NeRF and generally beneficial for low-quality or unknown camera poses (Lin et al., 2021; Fu et al., 2023; Chng et al., 2022; Lin et al., 2023; Meng et al., 2021; Bian et al., 2023). By optimizing the camera poses during the

training, the quality of the rendering and the densified point cloud from Gaussian Splatting could reach the quality of the SfM without additional pre-processing time. Furthermore, the real-time capability of the HoloLens offers the potential to insert data into Gaussian Splatting during the optimization, which is, with regard to SLAM approaches (Fink et al., 2023; Rosinol et al., 2023; Zhu et al., 2022), quite appealing.

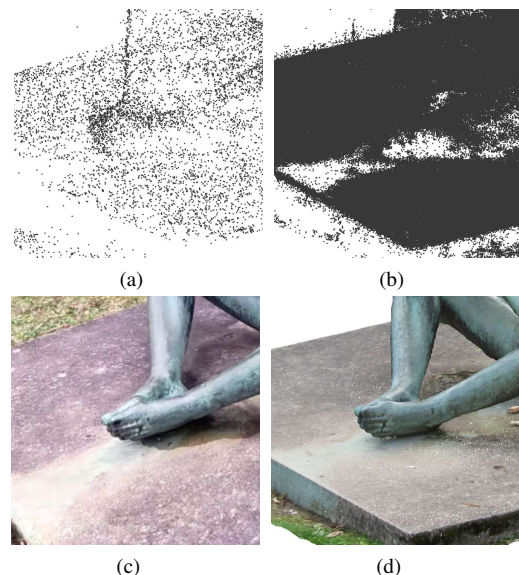


Figure 9. Low-textured, homogeneous surfaces. (a) Input SfM point cloud whose points are used as initial Gaussian centers. (b) Output densified point cloud of the Gaussian centers after training. (c) Rendered image. (d) Reference point cloud. It can be observed that the homogeneous, low-textured surfaces clearly have a lower point density of Gaussians.

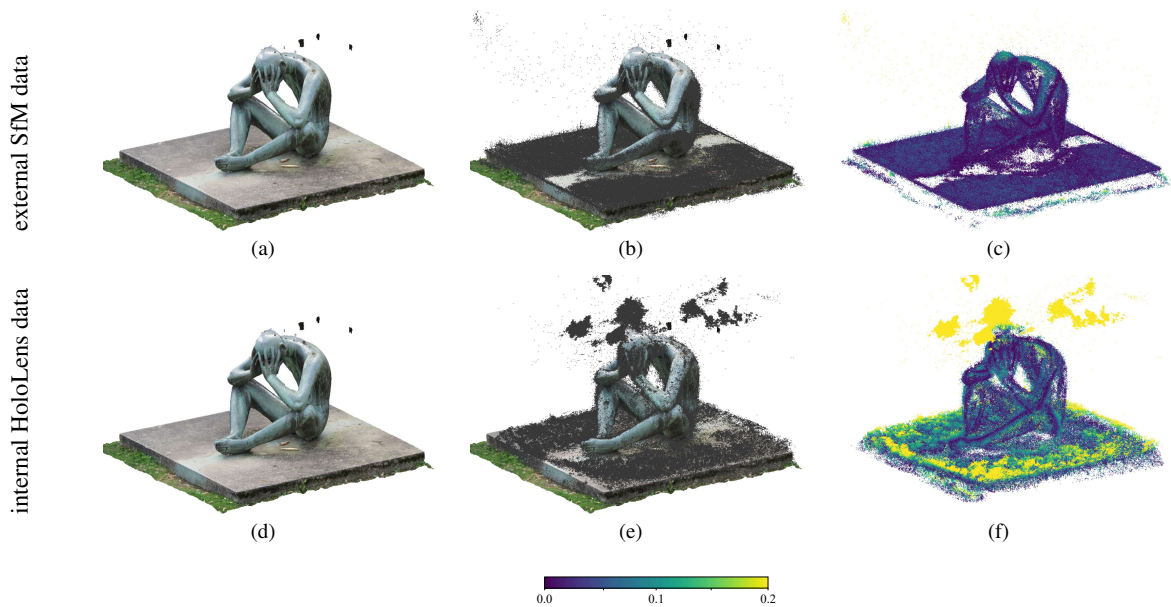


Figure 7. Geometric accuracy via Chamfer Distance \downarrow . Reference point cloud from MVS compared to the densified point clouds from Gaussian Splatting by extracting the center of each Gaussian. Top: Scene 'Denker' with external SfM data. From left to right: (a) Reference, (b) reference and GS point cloud, (c) Chamfer Distance of the GS point cloud. Bottom: Scene 'Denker' with internal HoloLens data. From left to right: (d) Reference, (e) reference and GS point cloud, (f) Chamfer Distance of the GS point cloud.

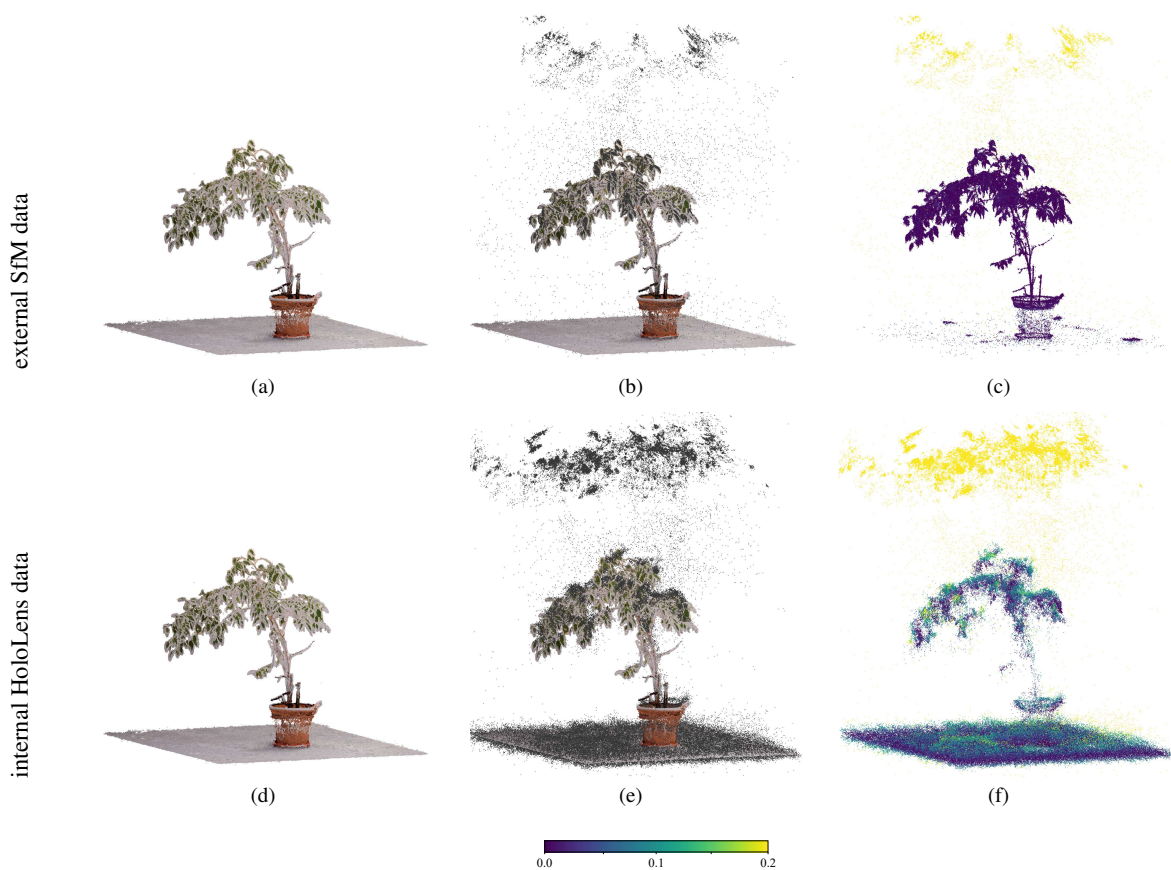


Figure 8. Geometric accuracy via Chamfer Distance \downarrow . Reference point cloud from MVS compared to the point clouds from Gaussian Splatting by extracting the center of each Gaussian. Top: Scene 'Ficus' with external SfM data. From left to right: (a) Reference, (b) reference and GS point cloud, (c) Chamfer Distance of the GS point cloud. Bottom: Scene 'Ficus' with internal HoloLens data. From left to right: (d) Reference, (e) reference and GS point cloud, (f) Chamfer Distance of the GS point cloud.

Generally, gaps in the densified point cloud, for both SfM and HoloLens data, are present, as shown in Figure 9. These gaps probably result from areas with homogeneous color, such as the areas around the legs of the statue in the scene 'Denker' and the low-textured pot of the plant in the scene 'Ficus'. Therefore, for the densified point cloud extraction, simply extracting the Gaussian centers is insufficient, due to the presence of floater artifacts and non-uniform point density on low-textured surfaces, where only a few individual Gaussians exist for homogeneous colors. In addition, it is not yet clear whether the Gaussian center or the Gaussian surface best represents the surface of the object geometry. These issues can be resolved by a suitable method for 3D point cloud extraction beyond querying of the Gaussian centers by further post-processing steps and extensions. In summary, despite the challenges, we see potential for HoloGS by combining the Microsoft HoloLens 2 with Gaussian Splatting for an instant 3D scene reconstruction and point cloud extraction.

6. Conclusion

In conclusion, HoloGS enables instant 3D scene reconstruction with 3D Gaussian Splatting directly from the internal sensor data of the Microsoft HoloLens 2. Although our results show some weaknesses compared to more elaborate approaches of using externally computed SfM data, such as a lower PSNR, floater artifacts and blurring in the rendered images as well as artifacts in the extracted point cloud, we nevertheless see potential for further optimization. In particular, refining the RGB camera poses during the training process could improve the results and enable real-time 3D reconstruction using state-of-the-art methods and entertainment devices like the HoloLens. As well as additional methods for point cloud and surface reconstruction from Gaussian Splatting. HoloGS thus represents a promising solution for using the Microsoft HoloLens 2 for instant 3D Gaussian Splatting, which offers further research potential in the realm of photogrammetry, computer vision and computer graphics.

ACKNOWLEDGEMENTS

The authors would like to thank Patrick Hübner from Technical University of Darmstadt for support regarding the HoloLens.

References

Bian, W., Wang, Z., Li, K., Bian, J.-W., Prisacariu, V. A., 2023. Nope-nerf: Optimising neural radiance field with no pose prior. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 4160–4169.

Chng, S.-F., Ramasinghe, S., Sherrah, J., Lucey, S., 2022. Gaussian activated neural radiance fields for high fidelity reconstruction and pose estimation. *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXIII*, Springer, 264–280.

Darmon, F., Basclé, B., Devaux, J.-C., Monasse, P., Aubry, M., 2022. Improving neural implicit surfaces geometry with patch warping. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6260–6269.

Dibene, J. C., Dunn, E., 2022. HoloLens 2 Sensor Streaming. *arXiv*. <https://arxiv.org/abs/2211.02648>.

Fink, L., Rückert, D., Franke, L., Keinert, J., Stamminger, M., 2023. Livenvs: Neural view synthesis on live rgb-d streams. *SIGGRAPH Asia 2023 Conference Papers*, 1–11.

Fu, H., Yu, X., Li, L., Zhang, L., 2023. Cbarf: Cascaded bundle-adjusting neural radiance fields from imperfect camera poses. *arXiv*: 2310.09776.

Haitz, D., Jutzi, B., Ulrich, M., Jäger, M., Hübner, P., 2023. Combining HoloLens with Instant-NeRFs: Advanced real-time 3D Mobile Mapping. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XLVIII-1/W1-2023, 167–174. <https://isprs-archives.copernicus.org/articles/XLVIII-1-W1-2023/167/2023/>.

Hou, J., Hübner, P., Schmidt, J., Iwaszczuk, D., 2024. Indoor Mapping with Entertainment Devices: Evaluating the Impact of Different Mapping Strategies for Microsoft HoloLens 2 and Apple iPhone 14 Pro. *Sensors*, 24(4).

Jäger, M., Hübner, P., Haitz, D., Jutzi, B., 2023. A Comparative Neural Radiance Field (NeRF) 3D Analysis of Camera Poses from HoloLens Trajectories and Structure from Motion. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XLVIII-1/W1-2023, 207–213. <https://isprs-archives.copernicus.org/articles/XLVIII-1-W1-2023/207/2023/>.

Jäger, M., Jutzi, B., 2023. 3D Density-Gradient based Edge Detection on Neural Radiance Fields (NeRFs) for Geometric Reconstruction. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XLVIII-1/W3-2023, 71–78. <https://isprs-archives.copernicus.org/articles/XLVIII-1-W3-2023/71/2023/>.

Jensen, R., Dahl, A., Vogiatzis, G., Tola, E., Aanaes, H., 2014. Large scale multi-view stereopsis evaluation. *2014 IEEE Conference on Computer Vision and Pattern Recognition*, IEEE, 406–413.

Jäger, M., Landgraf, S., Jutzi, B., 2023. Density uncertainty quantification with nerf-ensembles: Impact of data and scene constraints.

Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G., 2023. 3D Gaussian Splatting for Real-Time Radiance Field Rendering. *ACM Transactions on Graphics*, 42(4). <https://repo-sam.inria.fr/fungraph/3d-gaussian-splatting/>.

Li, Z., Müller, T., Evans, A., Taylor, R. H., Unberath, M., Liu, M.-Y., Lin, C.-H., 2023. Neuralangelo: High-fidelity neural surface reconstruction. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

Lin, C.-H., Ma, W.-C., Torralba, A., Lucey, S., 2021. Barf: Bundle-adjusting neural radiance fields. *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 5741–5751.

Lin, Y., Müller, T., Tremblay, J., Wen, B., Tyree, S., Evans, A., Vela, P. A., Birchfield, S., 2023. Parallel inversion of neural radiance fields for robust pose estimation. *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 9377–9384.

Lowe, D. G., 2004. Distinctive image features from scale-invariant keypoints. *International journal of computer vision*, 60, 91–110.

Meng, Q., Chen, A., Luo, H., Wu, M., Su, H., Xu, L., He, X., Yu, J., 2021. Gnerf: Gan-based neural radiance field without posed camera. *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 6351–6361.

Mildenhall, B., Srinivasan, P. P., Tancik, M., Barron, J. T., Ramamoorthi, R., Ng, R., 2020. NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis. *European Conference on Computer Vision (ECCV)*, 405–421.

Oechsle, M., Peng, S., Geiger, A., 2021. Unisurf: Unifying neural implicit surfaces and radiance fields for multi-view reconstruction. *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 5589–5599.

Park, J. J., Florence, P., Straub, J., Newcombe, R., Lovegrove, S., 2019. Deepsdf: Learning continuous signed distance functions for shape representation. *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 165–174.

Rosinol, A., Leonard, J. J., Carlone, L., 2023. Nerf-slam: Real-time dense monocular slam with neural radiance fields. *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, IEEE, 3437–3444.

Schönberger, J. L., Frahm, J.-M., 2016. Structure-from-motion revisited. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

Wang, P., Liu, L., Liu, Y., Theobalt, C., Komura, T., Wang, W., 2021. Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. *34*, 27171–27183.

Weinmann, M., Jäger, M. A., Wursthorn, S., Jutzi, B., Hübner, P., 2020. 3D indoor mapping with the Microsoft HoloLens: qualitative and quantitative evaluation by means of geometric features. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 1, 165–172.

Weinmann, M., Wursthorn, S., Weinmann, M., Hübner, P., 2021. Efficient 3D Mapping and Modelling of Indoor Scenes with the Microsoft HoloLens: A Survey. *Journal of photogrammetry, remote sensing and geoinformation science*, 89, 319–333.

Yariv, L., Gu, J., Kasten, Y., Lipman, Y., 2021. Volume rendering of neural implicit surfaces. *Advances in Neural Information Processing Systems*, 34, 4805–4815.

Zhang, J., Yao, Y., Li, S., Fang, T., McKinnon, D., Tsin, Y., Quan, L., 2022. Critical regularizations for neural surface reconstruction in the wild. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6270–6279.

Zhu, Z., Peng, S., Larsson, V., Xu, W., Bao, H., Cui, Z., Oswald, M. R., Pollefeys, M., 2022. Nice-slam: Neural implicit scalable encoding for slam. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12786–12796.