

# 3D Gaussian Splatting aided Localization for Large and Complex Indoor-Environments

Vincent Ress<sup>1</sup>, Jonas Meyer<sup>1,2</sup>, Wei Zhang<sup>1</sup>, David Skuddis<sup>1</sup>, Uwe Soergel<sup>1</sup>, Norbert Haala<sup>1</sup>

<sup>1</sup> Institute for Photogrammetry and Geoinformatics, University of Stuttgart, Germany - forename.lastname@ifp.uni-stuttgart.de

<sup>2</sup> Institute of Geomatics, University of Applied Sciences and Arts Northwestern Switzerland, Switzerland - jonas.meyer@fhnw.ch

**Keywords:** Dense Reconstruction, Indoor Localization, SLAM, Autonomous Indoor Construction

## Abstract

The field of visual localization has been researched for several decades and has meanwhile found many practical applications. Despite the strong progress in this field, there are still challenging situations in which established methods fail. We present an approach to significantly improve the accuracy and reliability of established visual localization methods by adding rendered images. In detail, we first use a modern visual SLAM approach that provides a 3D Gaussian Splatting (3DGS) based map to create reference data. We demonstrate that enriching reference data with images rendered from 3DGS at randomly sampled poses significantly improves the performance of both geometry-based visual localization and Scene Coordinate Regression (SCR) methods. Through comprehensive evaluation in a large industrial environment, we analyze the performance impact of incorporating these additional rendered views.

## 1. Introduction

Dense 3D scene reconstruction from imagery remains a core challenge in computer vision, robotics, and photogrammetry. Recent breakthroughs in deep learning, including 3D Gaussian Splatting (3DGS) (Kerbl et al., 2024), have significantly advanced both the performance and visual quality of the reconstruction. Within our work, we focus on 3D mapping of complex, large-scale indoor environments such as construction sites and factory halls. This initiative is driven by a project within the Cluster of Excellence Integrative Computational Design and Construction for Architecture (IntCDC) at the University of Stuttgart, which aims to enable autonomous indoor construction for new or preexisting buildings (IntCDC, 2024a).

Typical construction tasks, including material handling and element assembly, require highly accurate mapping approaches to enable precise localization of both building components and the construction robots. Image-based localization methods are particularly valuable due to the widespread availability and low cost of cameras, which are now standard equipment on most modern robots. This universal availability means localization algorithms can be deployed across diverse environments and robotic platforms with minimal modifications. In well-observed areas with rich texture, current image-based localization algorithms can achieve exceptional accuracy in determining both position and orientation. However, these methods' performance deteriorates significantly when confronted with substantial perspective differences or large baselines between reference data and images being processed for localization. Within the scope of this work, we present and evaluate a solution that is intended to reduce the influence of perspective differences in particular by increasing the number of reference data in synthesizing novel views.

An essential step in generating novel views is the collection of data and mapping of the environment. While Terrestrial Laser Scanning (TLS) has traditionally been the standard for high accurate 3D mapping in indoor and Global Navigation Satellite Systems (GNSS)-denied environments, it also comes with drawbacks - namely long processing and acquisition times, including the need for manual environment preparation. Mobile




Figure 1. Construction hall used as test scenario

mapping systems utilizing Simultaneous Localization and Mapping (SLAM)-based approaches have emerged as a more flexible and efficient alternative. However, despite offering real-time capabilities, visual SLAM algorithms typically produce less accurate and sparse maps, with particular vulnerability to scaling issues in large open spaces (Ress et al., 2024).

To support the project's long-term goal of enabling high-frequency, fully autonomous robotic surveys of construction sites, this study concentrates on efficient and adaptable SLAM-based acquisition algorithms. Additionally, a photogrammetric reference based on Structure-from-Motion (SfM) was created to benchmark the accuracy of image-based mapping algorithms. Using the computed camera poses and point clouds, a 3D Gaussian Splatting (3DGS) environment is initialized and optimized (cf. Chapter 2). Research by (Ress et al., 2024) has demon-

strated that this representation method is particularly effective in visualizing large and structurally complex indoor environments. The final stage demonstrates the viability of the captured 3D representation for localization purposes. The dense Gaussian Splatting environment is used to generate rendered RGB images, which then provide additional training data for the DSAC\* Network (Brachmann and Rother, 2021). This network architecture was selected for its state-of-the-art performance in RGB image localization, enabling scene coordinate estimation followed by a RANSAC-based camera pose computation for images to be localized. For comparison, we also implemented a geometry-based visual localization method that employs a global image descriptor to identify the most similar images with known poses, followed by a local bundle adjustment to query the pose of the input image (cf. Chapter 3).

## 2. Dense SLAM for 3D Gaussian mapping of a construction hall

As illustrated in the overview of our pipeline (Fig. 2), we introduce a dense visual SLAM pipeline for 3D Gaussian mapping of a construction hall, serving as the foundation for robust visual localization. This section is organized as follows: We first review related work in dense mapping and 3D scene representation (Sec. 2.1). We then describe our sensor system and data acquisition setup (Sec. 2.2), followed by details of our visual SLAM and 3D Gaussian optimization pipeline (Sec. 2.3). Next, we outline our photogrammetric reference to evaluate the accuracy of the visual SLAM pipeline (Sec. 2.4). Finally, we evaluate the quality of our Gaussian mapping through quantitative and qualitative comparisons (Sec. 2.5).

### 2.1 Related Work

Dense visual SLAM remains fundamental for autonomous navigation and 3D mapping in robotics and photogrammetry. While traditional SLAM systems like ORB-SLAM (Mur-Artal et al., 2015) use sparse feature matching for pose tracking, they often struggle in complex environments with low texture or repetitive patterns. Dense SLAM approaches such as LSD-SLAM (Engel et al., 2014) and DSO (Engel et al., 2017) address this by leveraging dense image information, improving robustness in challenging scenarios. More recently, learning-based systems like DROID-SLAM (Teed and Deng, 2021) have emerged, using neural networks to achieve more accurate pose estimation and complete 3D scene representations. However, the resulting point clouds and meshes have limitations - point clouds often lack fine details and textures, while high-resolution meshes require significant storage. 3D Gaussian Splatting (3DGS) (Kerbl et al., 2024) offers a promising alternative by providing a compact yet expressive scene representation that maintains both visual fidelity and geometric accuracy across viewpoints. By integrating Gaussian Splatting with dense SLAM, as demonstrated in HI-SLAM2 (Zhang et al., 2024a), we can produce detailed 3D representations of complex indoor environments like construction halls. This approach not only improves the quality of the reconstructed scene but also enables the rendering of novel views to enhance mapping coverage and robust visual localization.

### 2.2 Sensor System and Data Acquisition

A section of the Large-Scale Construction Robotics Laboratory (LCRL) within the Cluster of Excellence IntCDC was selected as the test environment (IntCDC, 2024). The facility comprises

a large construction hall that houses multiple robotic building part prefabrication plants, instructor workspaces, material storage areas, and traditional fabrication tools. The environment features both expansive open spaces and narrow corridors (cf. Fig. 1). Within the recorded area, various structures constructed from materials including concrete, steel, and wood are present.

The experimental setup utilizes a 6-wheeled robotic platform that provides essential functionalities including power supply, computational resources, and mobility capabilities (cf. Fig. 2). Data collection was performed using a camera head equipped with four ZED X RGB-D cameras. To optimize processing performance, only the forward-facing camera was utilized (1080p@10fps) and analyzed in the subsequent mapping procedures. To create a comprehensive reference map, the robot traversed all accessible areas within its physical dimensions. Where feasible, paths were covered in both directions. The localization data was collected during a separate, independent run to ensure translational and rotational deviations from the reference poses and allow for a thorough evaluation of localization performance. The total mapped area encompassed approximately  $2500m^2$ .

### 2.3 Visual SLAM

The goal of the Gaussian map is to provide a compact yet expressive scene representation that maintains high visual quality and geometric accuracy across different viewpoints as shown in Fig. 2. Inspired by (Ress et al., 2024), we first employ a dense visual SLAM (Teed and Deng, 2021) with stereo camera input to obtain camera poses and reconstruct a point cloud map of the observed scene, providing an initial sparse geometric structure of the environment. This point cloud serves as initialization for the 3D Gaussian primitives.

To create the final map, we optimize the parameters of the Gaussian primitives (position, orientation, scale, opacity, and color) through multiple loss functions. The primary objective is achieved through a photometric loss that measures color consistency between rendered and observed images. To enhance geometric accuracy (Zhang et al., 2024b), we incorporate stereo depth measurements as a geometric constraint, enforcing consistency between the rendered depths and the estimated stereo depths. Additionally, to handle varying exposures across different viewpoints (Kerbl et al., 2024, Zhang et al., 2024a), we estimate a  $3 \times 4$  affine transformation matrix for each training view to compensate for illumination changes. To further improve the map quality, we incorporate the anti-aliasing filter proposed by (Yu et al., 2024). In mobile robotics applications, varying observation distances result in changing sampling rates, which can lead to rendering artifacts. The filter effectively mitigates these by regularizing the size of Gaussian primitives. Through iterative optimization, these combined techniques guide the refinement of Gaussian primitives to achieve accurate geometric and appearance representation.

By leveraging our Gaussian map, we can generate highly detailed color and depth images from arbitrary viewpoints. This capability allows us to generate views even for areas that were not directly captured during data acquisition. Thus to enhance view coverage, we systematically augment our dataset by sampling 25 additional poses for each keyframe, varying both position and orientation. Along the robot's driving direction (longitudinal), we sample within  $\pm 0.5m$  translation, while allowing larger lateral translations of  $\pm 2.0m$  to expand coverage perpendicular to the path. For orientation, we sample across

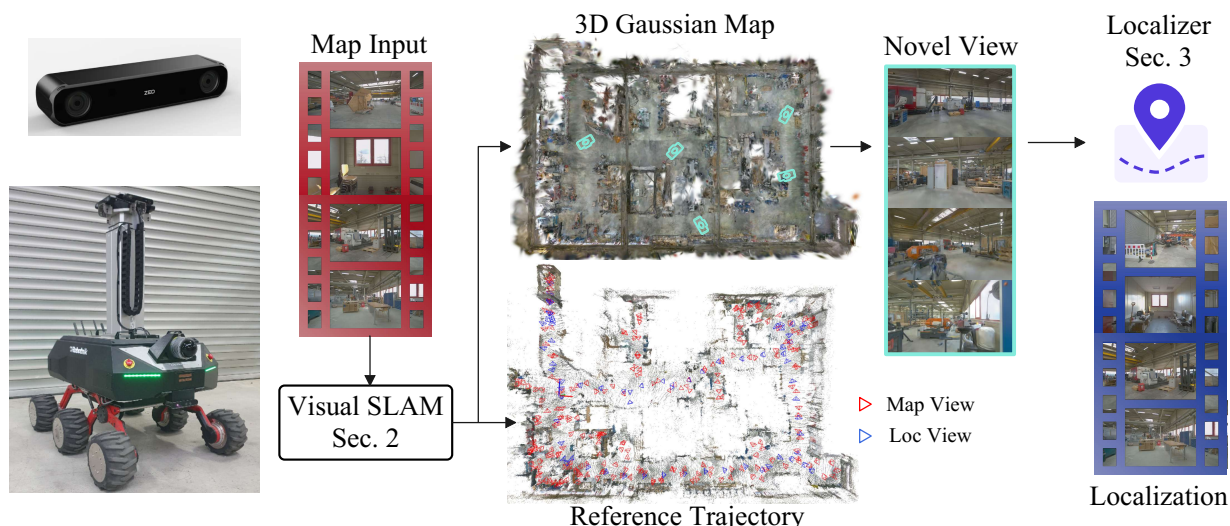


Figure 2. Overview of our pipeline: first stereo inputs are processed by dense visual SLAM to provide camera poses and initial scene structure, which are refined through 3D Gaussian Splatting to create a high-quality scene representation. This enables rendering of novel views to complement real images for robust visual localization in complex indoor environments.

the full  $\pm 180^\circ$  yaw rotation range to enable views from all directions. While these novel views from extrapolated poses may exhibit lower visual quality or sometimes significant artifacts compared to renderings at original viewpoints (Fig. 4), they provide valuable additional information for localization. As we demonstrate in Section 3.3, our localization approaches can effectively utilize these additional rendered views to improve localization performance.

## 2.4 Photogrammetric Reference

To evaluate the trajectory accuracy of our visual SLAM pipeline, we created a photogrammetric model with the SfM software Agisoft Metashape (Agisoft LLC, 2024). First, we jointly aligned all images of the mapping and the localization trajectory by using the image poses from the SLAM pipeline as pose priors. The test area lacks precise ground control points (GCPs) that could be used to minimize drift and scale issues. Therefore, we selected 26 distinctive points at strategically reasonable locations to mutually stabilize sub-loops of both trajectories. Each of these points was measured manually in around 20 to 50 images. In addition, we determined 10 reference scales distributed across the entire test area to ensure that the final model is true to scale. For the final camera optimization, the origin of the coordinate systems was set to align with the first image of the mapping trajectory, while all remaining image poses were neglected. The photogrammetric reference model has an Root Mean Square Error (RMSE) for the scale bars of  $0.003m$  and an RMSE for all re-projection errors of the manually measured tie points of  $0.652pix$ .

We evaluate the trajectory accuracy of our visual SLAM pipeline by comparing it with the photogrammetric reference model. Fig. 3 shows the trajectory of the visual SLAM pipeline color-coded by the Absolute Trajectory Error (ATE) relative to the photogrammetric reference. The ATE measures the Euclidean distance between corresponding camera positions after aligning both trajectories using a similarity transformation. Overall, the pipeline achieves an ATE RMSE of  $0.16m$ , with maximum deviations of  $0.49m$ . This demonstrates the high accuracy of our approach for indoor mapping applications while maintaining real-time performance.

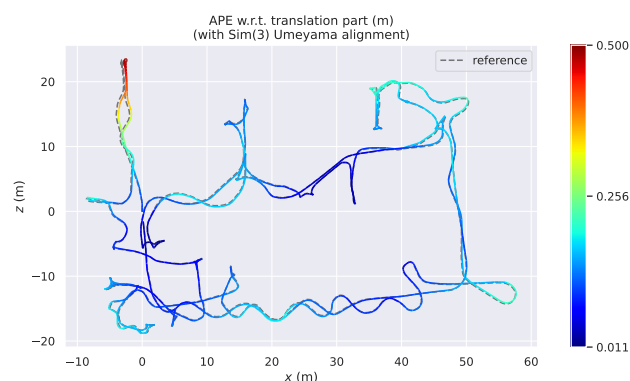


Figure 3. Absolute Trajectory Error of the visual SLAM pipeline relative to the photogrammetric reference.

## 2.5 Evaluation of Gaussian mapping quality

Config.	Description	PSNR
(a)	Color + depth Loss	26.14
(b)	(a) + exposure compensation	27.03
(c)	(b) + anti-aliasing filter	27.55

Table 1. Experiment configurations and map quality results.

We evaluate our Gaussian map quality through both quantitative metrics and qualitative visual comparisons between rendered and input images. Table 1 shows the systematic evaluation of three pipeline configurations using the Peak Signal-to-Noise Ratio (PSNR) metric. The baseline configuration using only color and depth loss achieves a PSNR of 26.14. Adding exposure compensation yields a significant improvement to 27.03 PSNR, while the final addition of the anti-aliasing filter further enhances the result to 27.55 PSNR.

As demonstrated in Fig. 4b), our method produces rendered views with high visual fidelity compared to the input images. The depth images show particularly notable improvements, with enhanced geometric detail preservation at object boundaries and better clarity of background structures at greater distances. Furthermore, our improved Gaussian map enables high-quality renderings even at extrapolated poses with large transla-



tional and rotational differences from the input views as shown in Fig. 4 d). This capability proves valuable for enhancing localization performance, as demonstrated in our experimental results in Section 3.2.

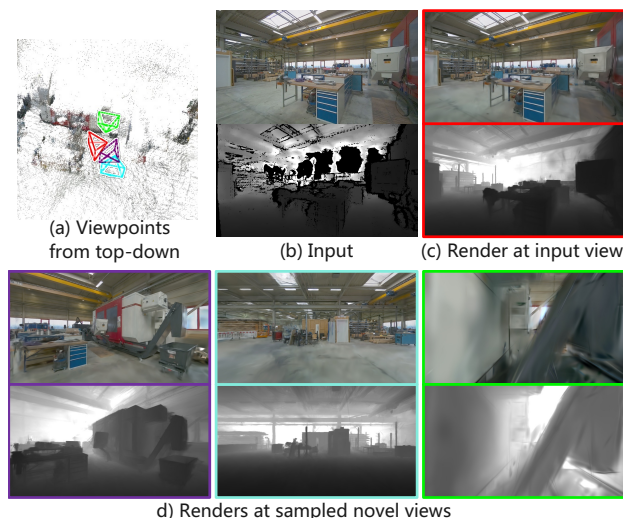


Figure 4. Qualitative results: (a) top-down view showing input viewpoint (red) and sampled novel viewpoints, (b) input image, (c) rendered image at input view, and (d) rendered images from sampled novel viewpoints.

### 3. Visual Localization

In this section, we present our visual localization approaches aiming at precise robot re-localization within the previously created reference map. This section is structured as follows: First, we provide a brief overview of related work (Sec. 3.1). Subsequently, we describe our employed approaches in detail (Sec. 3.2). Finally, we outline the conducted localization experiments and present the associated results (Sec. 3.3).

#### 3.1 Related Work

Visual (re-)localization describes the task of estimating a camera's pose with respect to a known scene. This task is a core component in different fields of research such as robotics, autonomous navigation and augmented reality. However, despite great efforts in research, visual localization is still facing major challenges when strong changes in viewpoint, illumination and appearance are present. We divide visual localization methods into two main categories: matching-based and regression-based methods. Regression-based methods can directly estimate a query image's pose based on the estimated scene coordinates, offering faster inference times compared to matching-based alternatives. In contrast, matching-based methods determine pose by actively comparing query images against a reference database. While this comparison process is more computationally intensive and the execution time depends on the database size, it currently achieves superior over all pose accuracy compared to regression-based approaches.

**3.1.1 Matching based Methods** Matching-based methods can be further distinguished between *image retrieval* and *geometry-based* methods.

Image retrieval methods aim at finding the  $k$  most similar images for a query image within a reference database. This is usually done by extracting descriptors and performing nearest

neighbor search in the descriptor domain. Traditionally, concepts such as Bag-of-Words (BoW) (Sivic and Zisserman, 2003), Histogram of Oriented Gradients (HOG) (Dalal and Triggs, 2005) or Vector of Locally Aggregated Descriptors (VLAD) (Jegou et al., 2010) were employed. With the advent of deep learning, learned global image descriptors have become prevalent. Approaches such as NetVLAD (Arandjelovic et al., 2016) or more recently EigenPlaces (Berton et al., 2023) and AnyLoc (Keetha et al., 2024) have advanced the state-of-the-art by using novel network architectures and training schemes. The accuracy of image retrieval methods is usually low, depending on the density of the reference images.

In contrast, geometry-based methods offer highly accurate pose estimates, in the range of a few centimeters. By employing correspondences between pixels in the query image and 3D points of the reference map, the query image's pose can be estimated via PnP/RANSAC (Gao et al., 2003, Fischler and Bolles, 1981). Over the last two decades an increasing amount of approaches to estimate those correspondences has been introduced. Starting from handcrafted local features such as SIFT (Lowe, 2004) in combination with nearest-neighbor matcher, over neural-based robust feature representations (Detone et al., 2018, Dushmanu et al., 2019, Zhao et al., 2023), novel feature matching methods (Sarlin et al., 2020, Lindenberger et al., 2023, Edstedt et al., 2024) towards pixel-wise dense matches (Sun et al., 2021, Chen et al., 2022). Additionally, hierarchical approaches leveraging image retrieval to obtain pose priors for search space reduction (Sarlin et al., 2019) as well as different map representations such as meshes (Panek et al., 2022), NeRFs (Zhou et al., 2025) and 3D Gaussian Splatting (Botashev et al., 2024, Zhai et al., 2024) have been proposed. All these approaches aim at estimating better correspondences from more challenging images or viewpoints and consequently, estimate the query image's pose more robustly and more accurately.

**3.1.2 Regression based Methods** Regression based methods for visual localization can be divided into Scene Coordinate Regression (SCR) and Direct Pose Regression (DPR).

In SCR 3D scene coordinates are estimated for all pixels of a given query image. While early approaches used regression trees to perform this task (Shotton et al., 2013) (Valentin et al., 2015), neural networks have become established in modern SCR approaches (Li et al., 2018) (Dong et al., 2022) (Wang et al., 2024). While most approaches require an extensive training phase before the inference phase, there are also some scene-agnostic approaches (Yang et al., 2019) (Revaud et al., 2024). The estimation of the scene coordinates is usually followed by a RANSAC (Fischler and Bolles, 1981) based pose estimation.

A series of approaches adds RANSAC-based pose estimation to the training of scene coordinates, achieving excellent results (Brachmann et al., 2017) (Brachmann and Rother, 2021). While here, despite the final training in an end-to-end fashion, is still a clear distinction between SCR and pose estimation in the internal structure, DPR methods use a single neural network to directly estimate poses for query images (Kendall et al., 2015) (Wang et al., 2020) (Wang et al., 2023).

#### 3.2 Localization Approaches

In order to, re-localize our robot within the reference map, we use two different approaches. A Hierarchical Geometry-based Visual Localization (HGVL) (cf. Sec. 3.2.1) approach and a Scene Coordinate Regression (SCR) (cf. Sec. 3.2.2) approach.



### 3.2.1 Hierarchical Geometry-based Visual Localization

Our hierarchical geometry-based visual localization pipeline has a comparable design as the one presented in (Meyer et al., 2020). First, the ten most similar images to the query image are selected via image retrieval. Subsequently, a local SfM model is built via triangulation of pair-wise matched images from known poses. Finally, the pose of the query image is estimated based on the established 2D-3D correspondences using a PnP solver and a RANSAC scheme.

For image retrieval we use EigenPlaces (Berton et al., 2023). Thanks to its novel training scheme, which incorporates images from various different and also strongly differing viewpoints, EigenPlaces convinces by increased viewpoint invariance than competing image retrieval approaches. This should lead to more robust and thus more favorable constellations regarding the SfM-based scene reconstruction step. Further, we use ALIKED (Zhao et al., 2023) a local feature detection and description method and LightGlue (Lindberger et al., 2023) as feature matching method. The combination of ALIKED and LightGlue has proven to be very efficient and performant when matching images under challenging conditions (Bellavia et al., 2024). Local scene reconstruction and registration of the query image is performed using COLMAP (Schönberger and Frahm, 2016).

In a preprocessing step, global image descriptors and local features are extracted for all reference images. Global image descriptors are stored in a database for fast nearest neighbor search in the descriptor domain. Similarly, local features are written to file to be swiftly read during the localization phase.

**3.2.2 Scene Coordinate Regression** Based on its state-of-the-art performance across benchmark datasets, including the 7 scenes (Shotton et al., 2013) and Cambridge dataset (Kendall et al., 2015), (Brachmann and Rother, 2021)'s DSAC\* network was selected to evaluate SCR-based approaches. The network's architecture consists of two main components: a Fully Convolutional Neural Network (FCNN) to estimate (3D) Scene Coordinates and a subsequent Differential RANSAC (DSAC) module for pose determination.

The initial training phase uses a loss function mainly based on the re-projection error. By minimizing the projection error the underlying information of the captured environment such as gray values and geometry are implicitly incorporated in the trained network. Despite increasing the training complexity, the input was restricted to RGB images (keyframes selected during the mapping process) and their corresponding poses to enhance flexibility and reduce dependency on specific sensor systems. The network's generalization capability was improved through data augmentation, specifically by varying brightness, contrast, and orientation. However, to optimize convergence behavior, augmentation was temporarily disabled during the initial training cycles. To optimize computational efficiency, the images for both inference and training were down-scaled to 25% of their original size, as we according to the authors (Brachmann and Rother, 2021) found this reduction caused no significant decrease in the localization performance.

In the final training step, the network directly minimizes the pose error computed from the estimated scene coordinates. The end-to-end differentiability of the DSAC method (incl. a gradient-based approximation of PnP algorithm), as established by its authors (Brachmann and Rother, 2021), enables iterative optimization throughout the training process.

### 3.3 Experimental Evaluation and Results

Approach	0.5°/ 0.02m	1.5°/ 0.05m	3°/ 0.1m	5°/ 0.5m	10°/ 1m
HGVL (a)	40.9	77.5	85.7	90.4	92.4
HGVL (b)	69.3	89.8	92.3	95.3	97.4
SCR (a)	0.0	0.0	1.5	31.6	56.4
SCR (b)	0.0	0.3	2.3	66.1	80.4

Table 2. Localization results obtained with both methods (HGVL / SCR) on reference data sets; (a) keyframes only, (b) keyframes and renderings from 3DGS; Localization results reported as the percentage of images localized within given translation and rotation thresholds.

To evaluate the accuracy of our visual localization approaches, we first establish reference poses by aligning the trajectory of localization data with the reference map using the dense alignment method from (Pan et al., 2024). This method achieves centimeter-level precision through global bundle adjustment, making it suitable for reference poses. We then compare the performance of our Hierarchical Geometry-based Visual Localization (HGVL) and Scene Coordinate Regression (SCR) approaches across different error thresholds. For each query image, we compute the absolute pose error as the deviation from the reference poses in both translation (meters) and rotation (degrees). The percentage of images successfully localized within various error thresholds is presented in Table 2. The results demonstrate that using rendered images from additional viewpoints/perspectives substantially improves localization accuracy for both methods evaluated. Notably, even in sparsely observed regions where rendered images exhibit significant blur to the point of being unrecognizable, this degradation does not appear to significantly impact the overall performance.

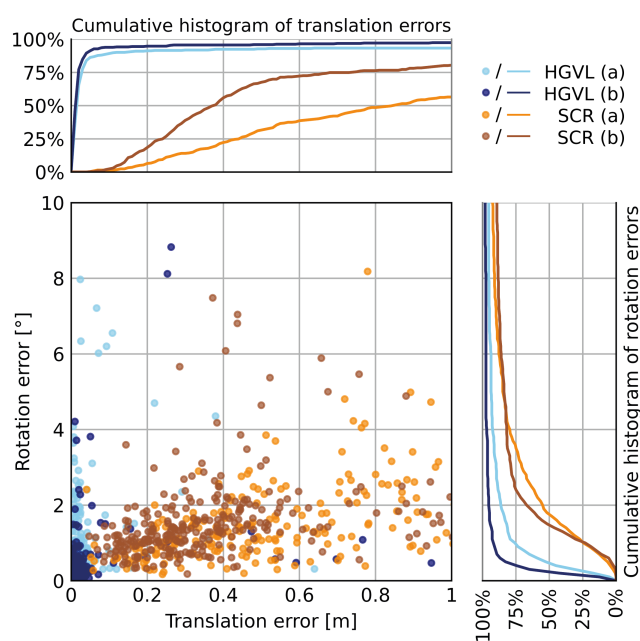


Figure 5. Localization results obtained with both methods (HGVL / SCR) on reference data sets (a) keyframes only (b) keyframes and renderings from 3DGS. Pose errors are separated into translation and rotation components.

This improvement is most pronounced in the SCR-based method, which shows a 24% increase in successful image localizations ( $< 10^\circ/1m$ ). The enhancement is evident in all accuracy categories. Analysis of Figure 5 and Figure 6 reveals that for the SCR-based method, apart from a small number of outliers, the addition of rendered training data leads to significant improvements in translation accuracy, while rotational accuracy shows only minimal change. We assume that the highly accurate poses and corresponding renderings of additional perspectives within the Gaussian environment enable a better capture of the geometry and thus a more accurate estimation of the scene coordinates. The enhanced geometric reconstruction particularly benefits depth estimation relative to the (virtual) camera, which explains the substantial reduction in translation error. Although incorporating rendered images doubled the required training time, the pose estimation inference time primarily depends on two factors: the network architecture and the number of RANSAC hypotheses. In our experiments, pose estimation for a single image required approx. 42ms.

The HGVL-based method exhibits different characteristics. Due to its already superior baseline performance, the increase in successfully localized images is more modest at 5%. However, there is a notable 29% improvement in high-precision localizations ( $< 0.5^\circ/0.02cm$ ), indicating a significant enhancement in the method's overall accuracy. Interestingly, the HGVL-based method shows particular improvement in rotational accuracy, while maintaining its already excellent translation accuracy. This behavior may be attributed to the hierarchical image selection strategy. The inclusion of rendered images increases the likelihood of finding reference views with higher overlap to the query image. This enhanced view selection stabilizes the subsequent local bundle adjustment, ultimately improving the query image alignment and the resulting pose. We assume that especially local bundles of keyframes representing pure forward motion sequences - typically challenging for SfM approaches - benefit greatly from the rendered images (cf. Figure 6). The execution times strongly depend on the number of reference images. In our setup with 29k rendered reference images, processing each query image took approximately 3-4 seconds.

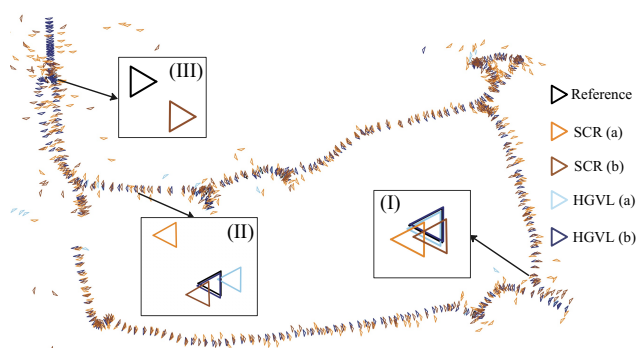


Figure 6. Qualitative localization results from SCR (a) in light brown to SCR (b) in dark brown and from HGVL (a) in light blue to HGVL (b) in dark blue; (a) keyframes only, (b) keyframes and renderings from 3DGS; Three detailed examples: (I) both approaches perform well, (II) HGVL outperforms SCR, and (III) only SCR (b) succeeds.

## 4. Conclusion

In this work, we demonstrate that modern visual SLAM methodology can achieve highly accurate environment mapping while enabling photorealistic rendering, primarily through the integration of 3D Gaussian Splatting. Our evaluation of both scene coordinate regression and geometry-based image localization approaches shows significant improvements through the introduction of rendered novel viewpoints. By incorporating rendered imagery into reference and training datasets, we achieve dual benefits: an increased number of successfully localized images and enhanced overall accuracy for both evaluated approaches.

Since construction tasks are typically pre-planned, with robotic platform working locations known days in advance. This planning horizon enables a practical scenario where additional imagery can be rendered specifically for these locations and integrated in the reference data. In this context, the training and pre-processing times required by our approaches become negligible. The rapid inference times of the Scene Coordinate Regression (SCR) Network enable real-time localization during navigation. More computationally intensive but highly accurate methods like our Hierarchical Geometry-based Visual Localization (HGVL) approach can then be employed for precise final localization prior to critical operations such as drilling or assembly. However, the SCR-based approach presented here requires enhanced accuracy for construction applications.

To address this limitation, we propose leveraging the superior depth information provided by 3D Gaussian Splatting during SCR training. This approach aims to achieve two objectives: improving network convergence and enhancing the accuracy of estimated scene coordinates. Furthermore, the network can provide an initial pose estimate for the query image, enabling filtered selection of reference images and subsequent local bundle adjustment for precise alignment.

## 5. Acknowledgements

Supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – EXC 2120/1 – 390831618.

## References

- Agisoft LLC, 2024. Agisoft Metashape Professional, Version 2.1.2.
- Arandjelovic, R., Gronat, P., Torii, A., Pajdla, T., Sivic, J., 2016. NetVLAD: CNN architecture for weakly supervised place recognition. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 5297–5307.
- Bellavia, F., Matas, J., Mishkin, D., Morelli, L., Remondino, F., Sun, W., Tabb, A., Trulls, E., Yi, K. M., Dane, S., Chow, A., 2024. Image matching challenge 2024 - hexathlon. <https://kaggle.com/competitions/image-matching-challenge-2024>. Kaggle.
- Berton, G., Trivigno, G., Caputo, B., Masone, C., 2023. EigenPlaces: Training Viewpoint Robust Models for Visual Place Recognition. *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 11046–11056.

- Botashev, K., Pyatov, V., Ferrer, G., Lefkimmiatis, S., 2024. Gslloc: Visual localization with 3d gaussian splatting. *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 5664–5671.
- Brachmann, E., Krull, A., Nowozin, S., Shotton, J., Michel, F., Gumhold, S., Rother, C., 2017. Dsac-differentiable ransac for camera localization. *Proceedings of the IEEE conference on computer vision and pattern recognition*, 6684–6692.
- Brachmann, E., Rother, C., 2021. Visual camera re-localization from RGB and RGB-D images using DSAC. *IEEE transactions on pattern analysis and machine intelligence*, 44(9), 5847–5865.
- Chen, H., Luo, Z., Zhou, L., Tian, Y., Zhen, M., Fang, T., McKinnon, D., Tsin, Y., Quan, L., 2022. ASpanFormer: Detector-Free Image Matching with Adaptive Span Transformer. S. Avidan, G. Brostow, M. Cissé, G. M. Farinella, T. Hassner (eds), *Computer Vision – ECCV 2022*, 13692, Springer Nature Switzerland, 20–36.
- Dalal, N., Triggs, B., 2005. Histograms of oriented gradients for human detection. *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, 886–893.
- Detone, D., Tomasz, M., Rabinovich, A., 2018. SuperPoint: Self-Supervised Interest Point Detection and Description. *Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*.
- Dong, S., Wang, S., Zhuang, Y., Kannala, J., Pollefeys, M., Chen, B., 2022. Visual localization via few-shot scene region classification. *2022 International Conference on 3D Vision (3DV)*, IEEE, 393–402.
- Dusmanu, M., Rocco, I., Pajdla, T., Pollefeys, M., Sivic, J., Torii, A., Sattler, T., 2019. D2-Net: A Trainable CNN for Joint Description and Detection of Local Features. *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 8084–8093.
- Edstedt, J., Sun, Q., Böckman, G., Wadenbäck, M., Felsberg, M., 2024. Roma: Robust dense feature matching. *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 19790–19800.
- Engel, J., Koltun, V., Cremers, D., 2017. Direct sparse odometry. *IEEE transactions on pattern analysis and machine intelligence*, 40(3), 611–625.
- Engel, J., Schöps, T., Cremers, D., 2014. Lsd-slam: Large-scale direct monocular slam. *European conference on computer vision*, Springer, 834–849.
- Fischler, M. A., Bolles, R. C., 1981. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6), 381–395.
- Gao, X.-S., Hou, X.-R., Tang, J., Cheng, H.-F., 2003. Complete solution classification for the perspective-three-point problem. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(8), 930–943.
- IntCDC, 2024. Intcdc large-scale construction robotics laboratory (lcr1). <https://www.intcdc.uni-stuttgart.de/research/research-infrastructure/>.
- Jegou, H., Douze, M., Schmid, C., Perez, P., 2010. Aggregating local descriptors into a compact image representation. *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 3304–3311.
- Keetha, N., Mishra, A., Karhade, J., Jatavallabhula, K. M., Scherer, S., Krishna, M., Garg, S., 2024. AnyLoc: Towards Universal Visual Place Recognition. *IEEE Robotics and Automation Letters*, 9(2), 1286–1293.
- Kendall, A., Grimes, M., Cipolla, R., 2015. PoseNet: A convolutional network for real-time 6-dof camera relocalization. *Proceedings of the IEEE international conference on computer vision*, 2938–2946.
- Kerbl, B., Meuleman, A., Kopanas, G., Wimmer, M., Lanvin, A., Drettakis, G., 2024. A hierarchical 3d gaussian representation for real-time rendering of very large datasets. *ACM Transactions on Graphics (TOG)*, 43(4), 1–15.
- Li, X., Ylioinas, J., Kannala, J., 2018. Full-frame scene coordinate regression for image-based localization. *arXiv preprint arXiv:1802.03237*.
- Lindenberg, P., Sarlin, P.-E., Pollefeys, M., 2023. LightGlue: Local Feature Matching at Light Speed. *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, IEEE, 17581–17592.
- Lowe, D. G., 2004. Distinctive Image Features from Scale-Invariant Keypoints. *International Journal of Computer Vision*, 60(2), 91–110.
- Meyer, J., Rettenmund, D., Nebiker, S., 2020. Long-Term visual Localization in Large Scale Urban Environments Exploiting Street Level Imagery. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 5(2), 57–63.
- Mur-Artal, R., Montiel, J. M. M., Tardos, J. D., 2015. ORB-SLAM: a versatile and accurate monocular SLAM system. *IEEE transactions on robotics*, 31(5), 1147–1163.
- Pan, P., Zhang, W., Haala, N., 2024. Dense collaborative mapping with deep visual SLAM method. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XLVIII-1-2024, 547–554.
- Panek, V., Kukulova, Z., Sattler, T., 2022. Meshloc: Mesh-based visual localization. S. Avidan, G. Brostow, M. Cissé, G. M. Farinella, T. Hassner (eds), *Computer Vision – ECCV 2022*, Springer Nature Switzerland, Cham, 589–609.
- Ress, V., Zhang, W., Skuddis, D., Haala, N., Soergel, U., 2024. SLAM for Indoor Mapping of Wide Area Construction Environments. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, X-2-2024, 209–216. <https://isprs-annals.copernicus.org/articles/X-2-2024/209/2024/>.
- Revaud, J., Cabon, Y., Brégier, R., Lee, J., Weinzaepfel, P., 2024. Sacreg: Scene-agnostic coordinate regression for visual localization. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 688–698.
- Sarlin, P.-E., Cadena, C., Siegwart, R., Dymczyk, M., 2019. From coarse to fine: Robust hierarchical localization at large scale. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 12708–12717.



- Sarlin, P.-E., DeTone, D., Malisiewicz, T., Rabinovich, A., 2020. SuperGlue: Learning Feature Matching With Graph Neural Networks. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, IEEE, 4937–4946.
- Schönberger, J. L., Frahm, J.-M., 2016. Structure-from-motion revisited. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 4104–4113.
- Shotton, J., Glocker, B., Zach, C., Izadi, S., Criminisi, A., Fitzgibbon, A., 2013. Scene coordinate regression forests for camera relocalization in rgb-d images. *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2930–2937.
- Sivic, Zisserman, 2003. Video Google: A text retrieval approach to object matching in videos. *Proceedings Ninth IEEE International Conference on Computer Vision*, 1470–1477 vol.2.
- Sun, J., Shen, Z., Wang, Y., Bao, H., Zhou, X., 2021. LoFTR: Detector-Free Local Feature Matching with Transformers. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, IEEE, 8918–8927.
- Teed, Z., Deng, J., 2021. Droid-slam: Deep visual slam for monocular, stereo, and rgb-d cameras. *Advances in neural information processing systems*, 34, 16558–16569.
- Valentin, J., Nießner, M., Shotton, J., Fitzgibbon, A., Izadi, S., Torr, P. H., 2015. Exploiting uncertainty in regression forests for accurate camera relocalization. *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4400–4408.
- Wang, B., Chen, C., Lu, C. X., Zhao, P., Trigoni, N., Markham, A., 2020. Atloc: Attention guided camera localization. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34number 06, 10393–10401.
- Wang, S., Kang, Q., She, R., Tay, W. P., Hartmannsgruber, A., Navarro, D. N., 2023. Robustloc: Robust camera pose regression in challenging driving environments. *Proceedings of the AAAI Conference on Artificial Intelligence*, 37number 5, 6209–6216.
- Wang, S., Laskar, Z., Melekhov, I., Li, X., Zhao, Y., Tolias, G., Kannala, J., 2024. Hscnet++: Hierarchical scene coordinate classification and regression for visual localization with transformer. *International Journal of Computer Vision*, 1–21.
- Yang, L., Bai, Z., Tang, C., Li, H., Furukawa, Y., Tan, P., 2019. Sanet: Scene agnostic network for camera localization. *Proceedings of the IEEE/CVF international conference on computer vision*, 42–51.
- Yu, Z., Chen, A., Huang, B., Sattler, T., Geiger, A., 2024. Mip-splatting: Alias-free 3d gaussian splatting. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 19447–19456.
- Zhai, H., Zhang, X., Boming, Z., Li, H., He, Y., Cui, Z., Bao, H., Zhang, G., 2024. SplatLoc: 3D Gaussian Splatting-based Visual Localization for Augmented Reality. *arXiv preprint arXiv:2409.14067*.
- Zhang, W., Cheng, Q., Skuddis, D., Zeller, N., Cremers, D., Haala, N., 2024a. HI-SLAM2: Geometry-Aware Gaussian SLAM for Fast Monocular Scene Reconstruction. *arXiv preprint arXiv:2411.17982*.
- Zhang, W., Sun, T., Wang, S., Cheng, Q., Haala, N., 2024b. HI-SLAM: Monocular Real-Time Dense Mapping With Hybrid Implicit Fields. *IEEE Robotics and Automation Letters*, 9(2), 1548–1555.
- Zhao, X., Wu, X., Chen, W., Chen, P. C. Y., Xu, Q., Li, Z., 2023. ALIKED: A Lighter Keypoint and Descriptor Extraction Network via Deformable Transformation. *IEEE Transactions on Instrumentation and Measurement*, 72, 1–16.
- Zhou, Q., Maximov, M., Litany, O., Leal-Taixé, L., 2025. The perfect match: Exploring nerf features for visual localization. A. Leonardis, E. Ricci, S. Roth, O. Russakovsky, T. Sattler, G. Varol (eds), *Computer Vision – ECCV 2024*, Springer Nature Switzerland, Cham, 108–127.