PlanGS: Active 3D Gaussian Reconstruction with Real-Time Planning

Supplementary Material

6. More Implementation Details

Parameter Setting. To determine whether a viewpoint is active (Sec. 3.3), we set $\sigma_u = 0.001$ and $\sigma_n = 0.1\%$. In Eq. 5, we use $\sigma_o = 0.5$ and $\sigma_d = 0.05$. In Eq. 7, the loss weights are $\lambda_d = 0.5$ and $\lambda_b = 0.01$. In Eq. 10, we set $\lambda_u = 0.5$. In Sec. 3.1, we set $\lambda_\gamma = 0.1$ in the depth-dependent exponential factor. In Eq. 13, we set $\sigma_p = 3$, and the heuristic weight λ_h is set to 0.4 for Replica and Gibson, and 0.7 for HM3D.

Poisson-disk Sampling Details. For Poisson-disk sampling, we first extract the obstacle points lying on the motion plane from the depth image within the camera's field of view. These obstacle points are then connected with the camera position to form a polygonal free-space boundary, within which Poisson-disk sampling is performed. To ensure that the sampled viewpoints are safe, we enforce a safety margin that keeps each sampled point at a sufficient distance from obstacles. When constructing the roadmap, two sampled points are connected only if the line segment between them lies entirely within the polygon, ensuring collision-free edges.

Real-world Experiments. When it comes to the deployment of the active reconstruction algorithm on a real-world robot, accurate pose estimation becomes a critical factor that limits reconstruction quality. To address this issue and ensure reliable, real-time pose acquisition over long-duration explorations, we employ FAST-LIO2 [47] algorithm using MID360 LiDAR devices integrated with an internal IMU. Compared to depth cameras, the LiDAR provides more accurate range measurements and, with its 360° field of view, can effectively handle severe geometric degeneracies in most indoor scenes.

Since the RGB sensor is not synchronized with the LiDAR device, we simply keep FAST-LIO2 running throughout the experiments. Whenever a pose is requested by PlanGS, we temporarily lock the quadruped robot to maintain a static state until the entire acquisition process is completed. During each acquisition, we obtain the current pose from FAST-LIO2 together with the corresponding RGB images. FAST-LIO2 outputs only the pose of the LiDAR device. To obtain the poses corresponding to the RGB images, the transformation between the LiDAR and the camera is required. Assuming a rigid relative configuration between the LiDAR and the camera, we perform the calibration prior to running the algorithm. The transformation matrix is obtained using the Direct Visual–LiDAR Calibration tool [25].

In order to show that our algorithm is capable of handling the exploration of many rooms, we build a barrier to mimic two-room scene. For details of the experiment, please refer to our video result.

(a) Per-step planner search time (b) Per-step roadmap update time

Figure 9. **Per-step planning time decomposition on Replica office 0.** We plot the roadmap update time and planner search time at each step.

Table 4. **Quantitative results of our method on the Replica dataset.**

Metric	Of0	Of1	Of2	Of3	Of4	R0	R1	R2	Avg.
PSNR↑	37.36	35.93	34.34	34.25	34.45	33.48	33.52	34.26	34.70
SSIM↑	0.96	0.95	0.95	0.94	0.94	0.93	0.92	0.94	0.94
LPIPS↓	0.10	0.15	0.11	0.13	0.15	0.16	0.17	0.15	0.14
Comp%↑	98.20	97.02	98.25	96.50	96.66	98.24	98.75	98.71	97.79
P.L.(m)↓	49.46	22.17	69.70	76.33	54.72	69.28	48.71	56.74	55.89
T_p(ms)↓	19.83	21.21	25.12	31.23	44.03	30.93	24.05	29.15	28.19
T_{all}(min)↓	2.23	0.83	4.10	4.74	2.56	5.05	2.72	2.78	3.12

Table 5. **Quantitative results of our method on the Gibson dataset.**

Metric	Beach	Denmark	Elmira	Eudora	Maida	Pablo	Ribera	Sumas	Avg.
PSNR↑	32.86	32.11	30.45	29.51	31.41	31.86	33.31	31.78	31.66
SSIM↑	0.93	0.93	0.89	0.90	0.91	0.92	0.93	0.92	0.92
LPIPS↓	0.16	0.16	0.23	0.26	0.21	0.21	0.17	0.20	0.20
Comp%↑	99.02	98.84	99.13	98.19	96.49	95.57	99.05	97.85	98.02
P.L.(m)↓	75.65	82.73	81.93	65.57	69.28	65.08	100.55	65.62	75.80
T_p(ms)↓	32.89	32.03	38.04	34.83	24.14	33.21	35.64	44.03	34.35
T_{all}(min)↓	4.12	4.61	5.24	3.27	3.77	3.09	7.44	3.72	4.41

7. More Runtime Analysis

We define the per-step planning time T_p as the total time required to (i) determine the next viewpoint (what to observe next), and (ii) compute the path from the agent current position to that viewpoint (how to reach it). In our implementation, T_p consists of two components: the roadmap update time and the planner search time.

As shown in Figure 9, we visualize these two components and T_p at each step on Replica office0 data. At the early stage of exploration, the scene contains large unknown

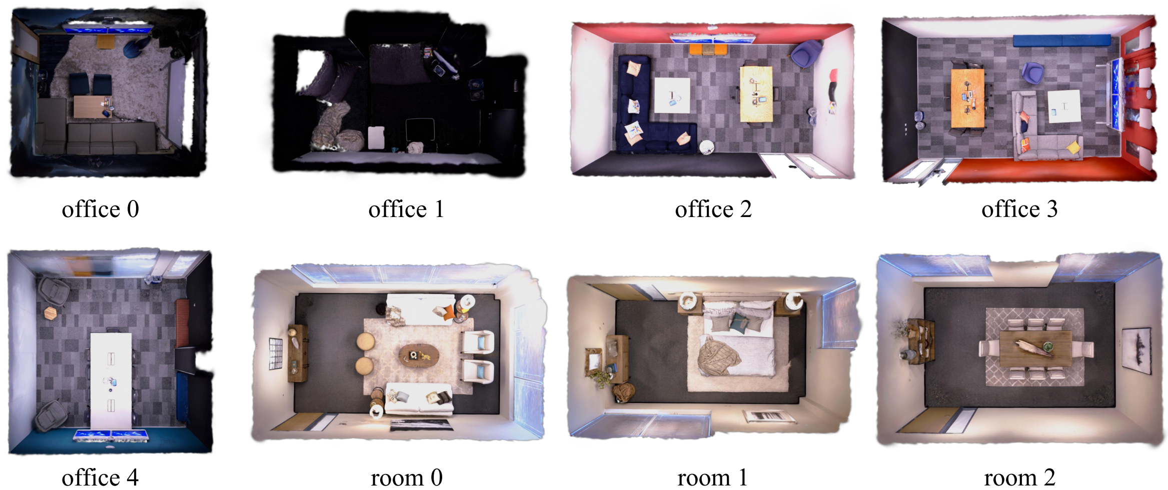

Figure 10. **Top-view visualization of our method on the Replica dataset.**

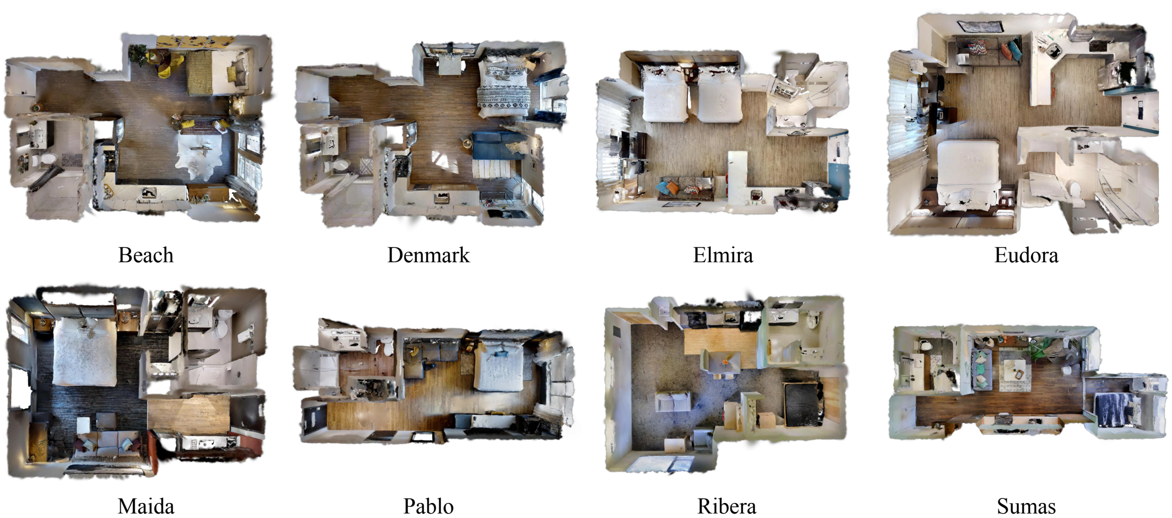

Figure 11. **Top-view visualization of our method on the Gibson dataset.**

regions; thus the roadmap undergoes substantial updates, leading to higher update time, while the planner search remains fast because most viewpoints are informative (active viewpoints). As exploration progresses, the roadmap stabilizes and its update time significantly decreases. In addition, the local region becomes well reconstructed, causing the planner to search farther in the roadmap to find informative viewpoints, which increases the planner search time.

Table 6. **Quantitative results of our method on the HM3D dataset.**

Metric	KJxdMPgweZG	nzuiinFMXvf	CFVBbU9Rsyb	mDdyQ6azhVD	Avg.
PSNR↑	30.13	31.33	29.53	31.07	30.52
SSIM↑	0.83	0.90	0.88	0.89	0.87
LPIPS↓	0.36	0.22	0.29	0.29	0.29
Comp%↑	96.51	97.47	97.45	96.68	97.03
P.L.(m)↓	153.89	187.35	177.98	123.64	160.71
T_p(ms)↓	59.00	57.42	62.24	61.84	60.13
T_{all}(min)↓	11.90	17.25	15.90	9.95	13.75

8. More Results

We present detailed quantitative results of our method on the Replica, Gibson, and HM3D datasets in Table 4, Table 5, and Table 6, reporting thorough metrics about mapping and planning. For mapping metrics, we report PSNR,

Figure 12. **Top-view visualization of our method on the HM3D dataset.**

SSIM, LPIPS and complete ratio metrics. As for planning metrics, we report path length, per-step planning time, and the overall runtime of active reconstruction. In addition to the top-view visualizations shown in Figure 10, Figure 11, and Figure 12, we invite readers to refer to our video results for clearer and more intuitive demonstrations.