AutoNeRF: Training Implicit Scene Representations with Autonomous Agents

Pierre Marza[1], Laëtitia Matignon[2], Olivier Simonin[1], Dhruv Batra[3, 5], Christian Wolf[4], Devendra Singh Chaplot[3]

INSA Lyon[1] UCBL[2] Meta AI[3] Naver Labs Europe[5] Georgia Tech[5]

ICCV23 PARIS
NeRF4ADR Workshop
Project Page

Introduction

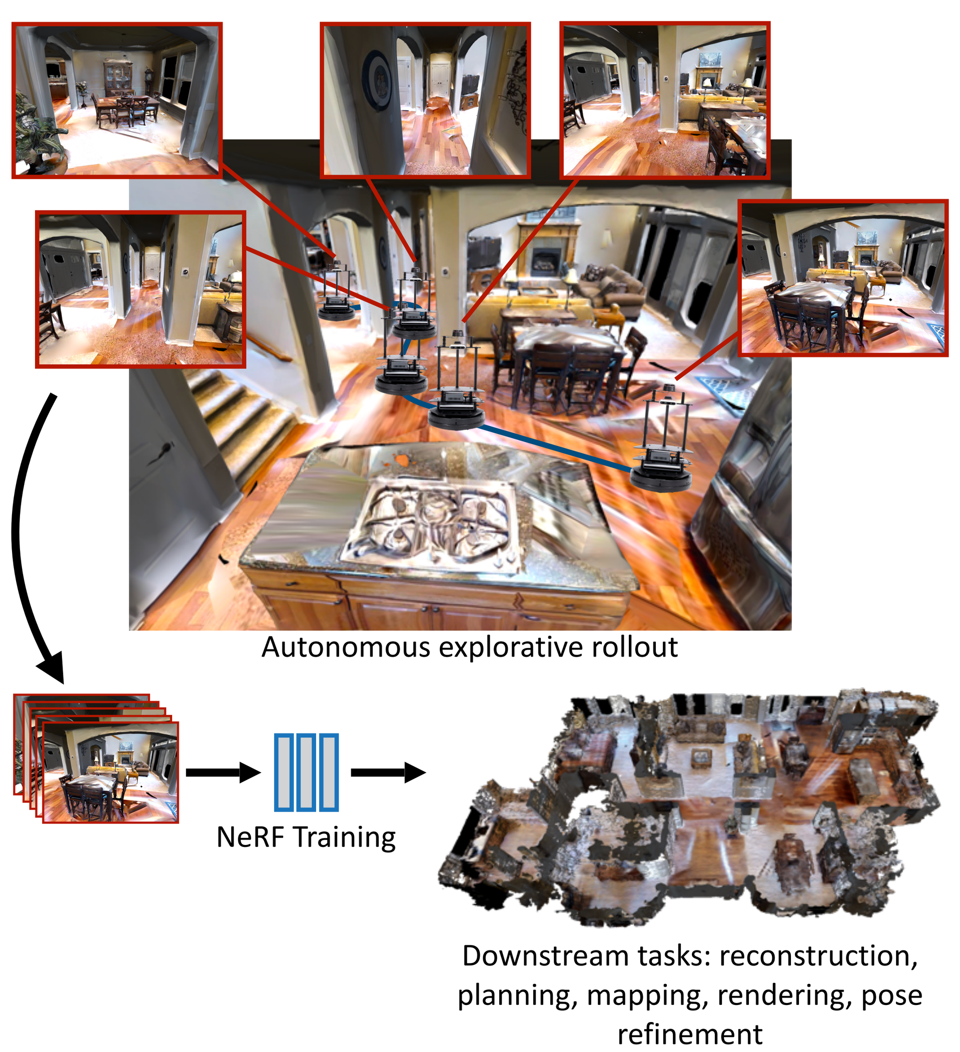

Contributions

- We introduce AutoNeRF, a modular policy trained with Reinforcement Learning (RL) that can explore an unseen 3D scene to collect data for training a NeRF model autonomously.
- While most prior work evaluates NeRFs on rendering quality, we propose a range of downstream tasks to evaluate them for Embodied AI applications.
- We show that AutoNeRF outperforms the well-known frontier exploration algorithm as well as learnt end-to-end counterparts, and we also study the impact of different RL reward functions on the downstream performance of the NeRF model.

Modular Policy

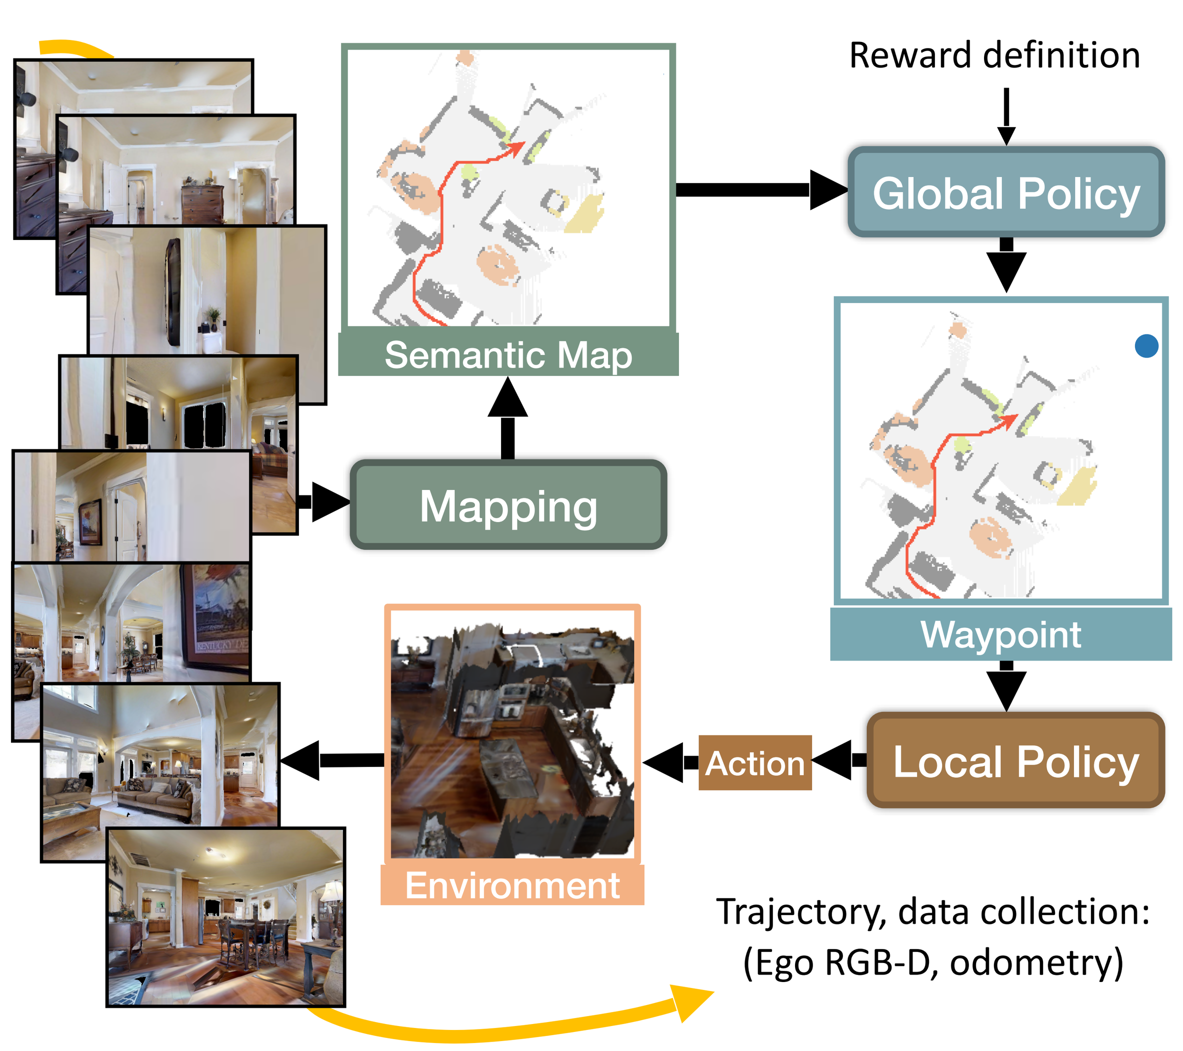

Downstream tasks

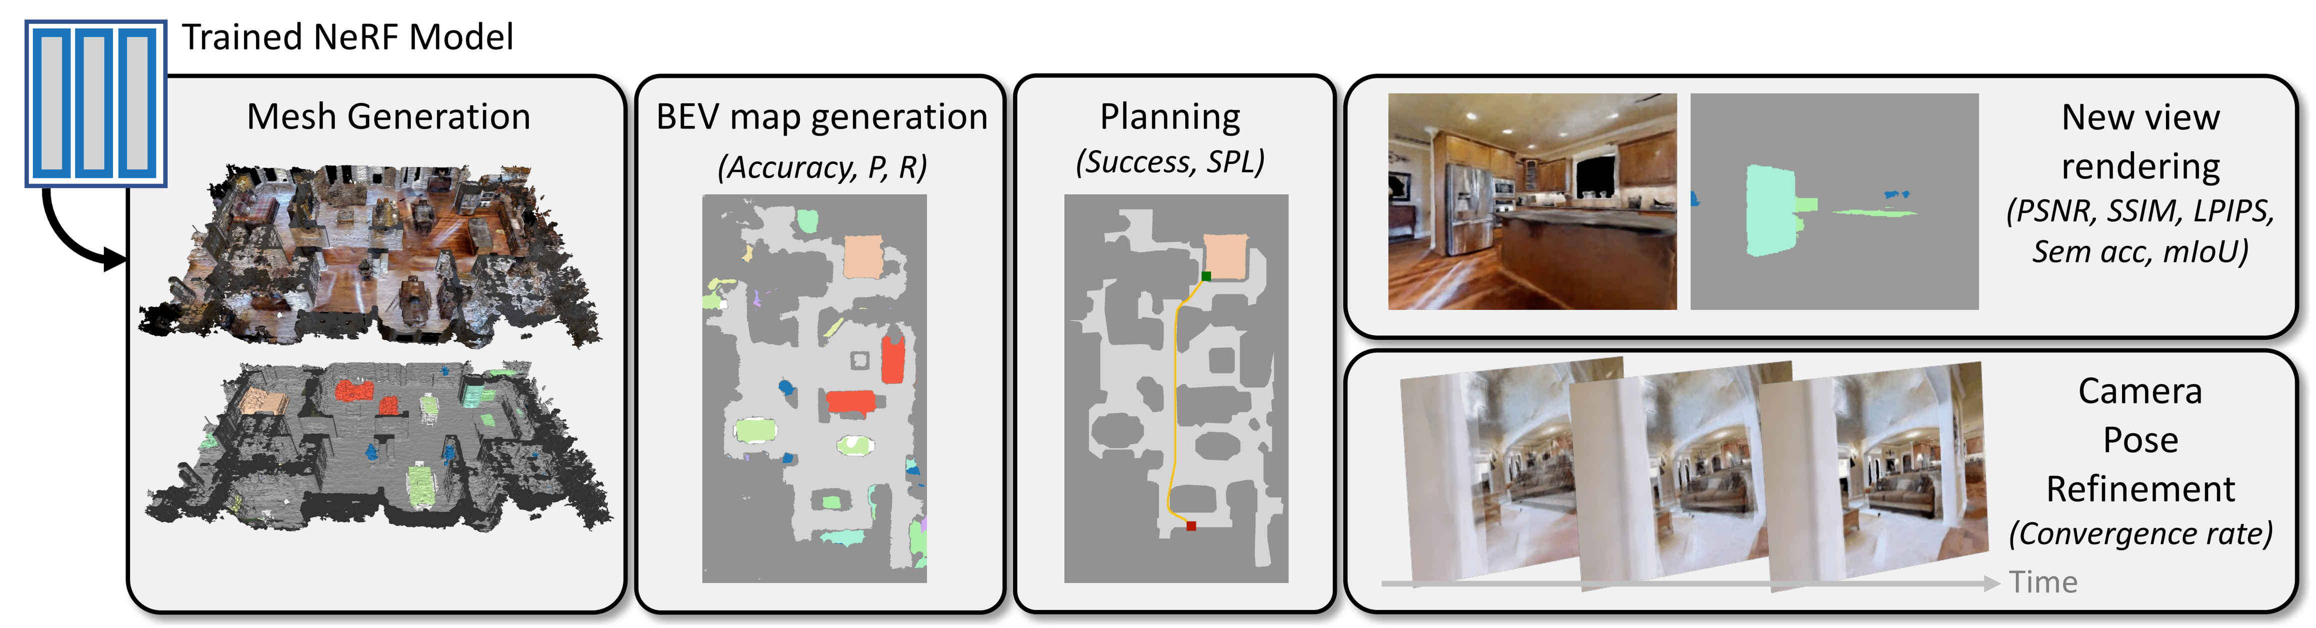

Mesh reconstruction

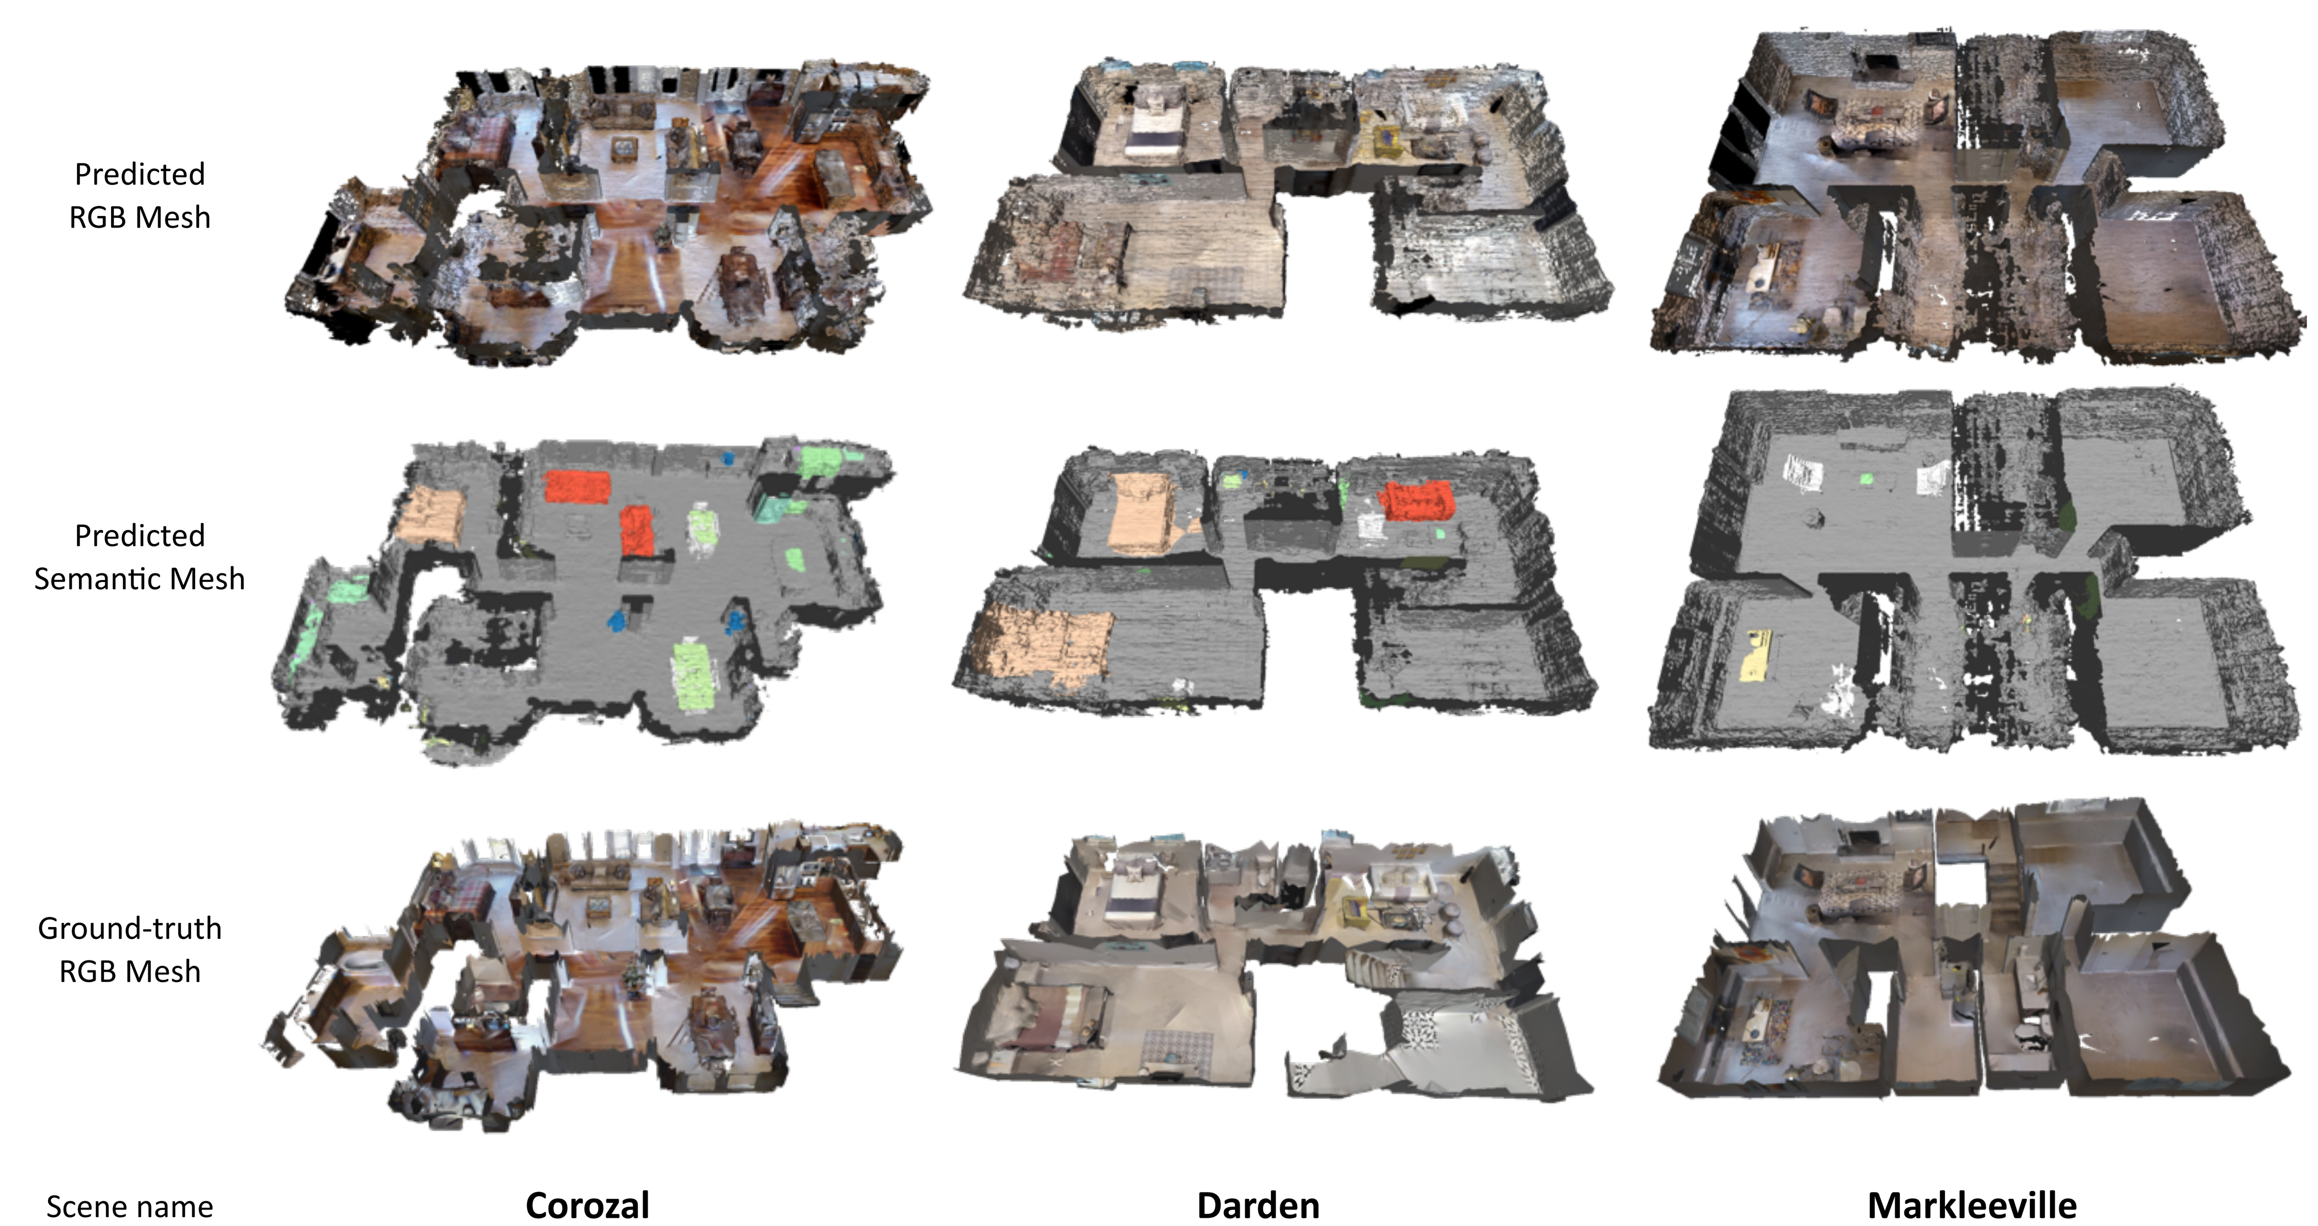

Qualitative results – BEV Maps

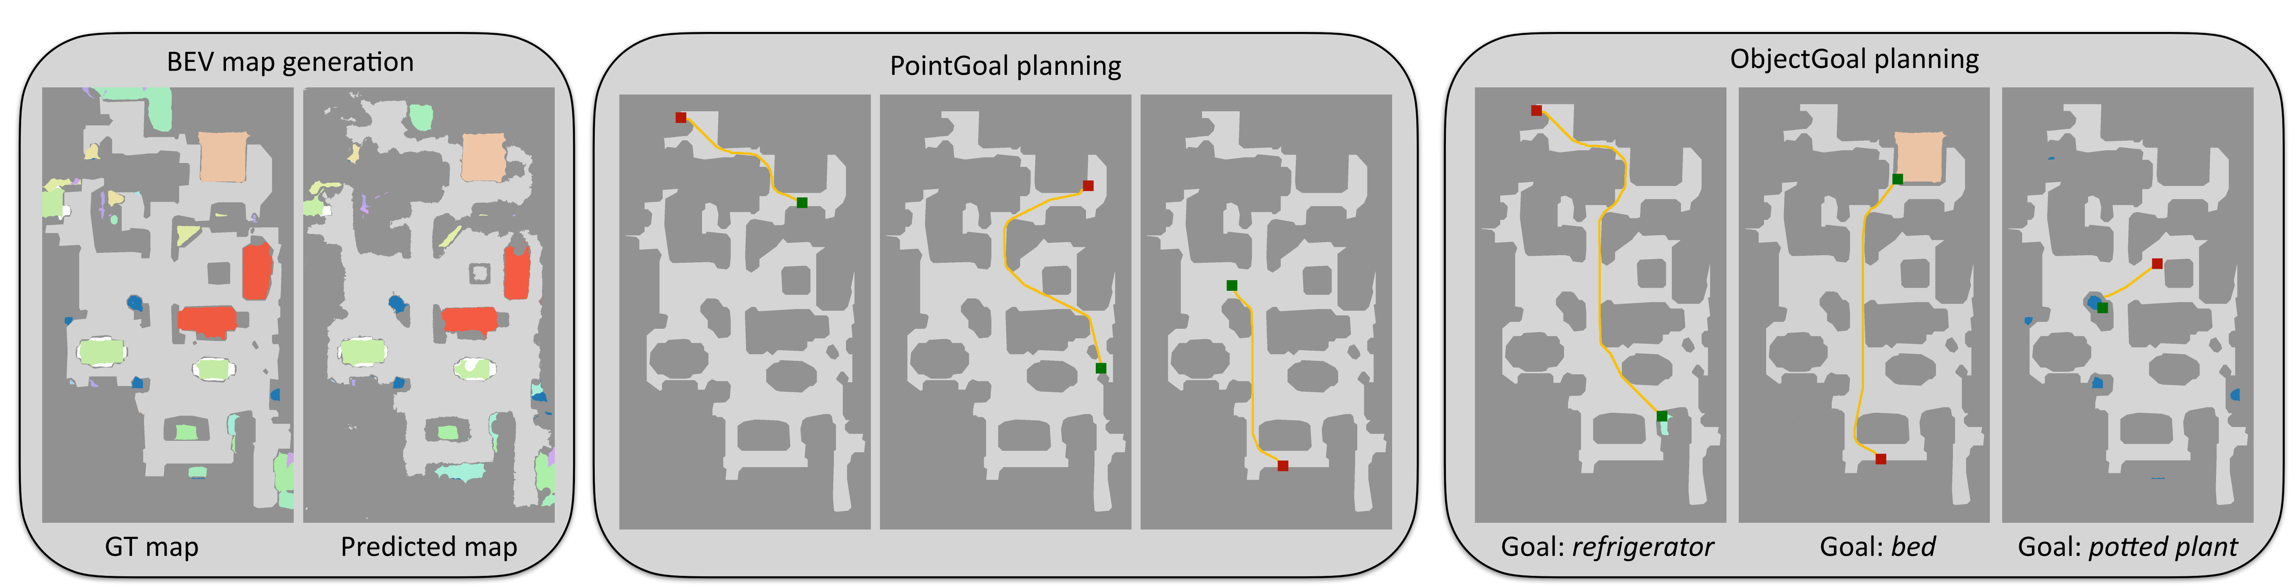

Qualitative results – Rendering

Autonomous adaptation to a new scene

An autonomous agent will likely struggle with specificities of a new environment. A safe solution is **scene-specific adaptation**: AutoNeRF is used to explore a scene and build a 3D representation that is loaded into a simulator to finetune a policy of interest, e.g. a depth-only PointGoal agent.

Policy	Success	SPL
Finetuned on Gibson meshes (not comparable)	99.7	97.9
Pre-trained (no finetuning)	90.2	82.9
Finetuned on AutoNeRF meshes	**92.9**	**86.7**

Quantitative results

Rewards – Maximizing...

Ours (cov.): Explored area
Ours (obs.): Seen Obstacles
Ours (sem.): Obs. in sem. map
Ours (view.): Viewpoints from which objects are seen

External baselines

Frontier: Frontier-based expl.
E2E: End-to-end RL agents
(*Ramakrishnan et al., An Exploration of Embodied Visual Exploration, IJCV 2021*)

	Rendering					Map Estimation					
	RGB			Semantics		Occupancy			Semantics		
Policy	PSNR	SSIM	LPIPS	Acc.	mIoU	Acc.	Prec.	Rec.	Acc	Prec.	Rec.
Frontier	19.75	0.743	0.343	81.4	65.7	81.2	86.9	49.9	99.7	26.6	21.0
E2E (cov.)	20.94	0.750	0.332	80.1	63.9	77.1	86.2	50.4	99.7	22.1	16.1
E2E (cur.)	20.60	0.747	0.338	78.7	61.9	81.8	90.3	50.7	99.7	19.2	12.5
E2E (nov.)	23.36	0.801	0.268	84.6	71.4	83.1	88.7	61.3	99.7	25.5	18.3
E2E (rec.)	23.17	0.797	0.270	84.1	70.5	81.6	87.6	60.0	99.7	26.2	18.0
Ours (cov.)	24.89	0.837	0.218	90.2	81.2	86.8	89.1	74.7	99.8	35.1	27.1
Ours (sem.)	25.34	0.843	0.207	**91.9**	81.8	86.6	88.3	76.5	99.8	35.7	29.8
Ours (obs.)	**25.56**	**0.846**	**0.203**	**91.8**	**83.2**	86.4	89.4	76.5	99.8	36.2	29.8
Ours (view.)	25.17	0.842	0.211	91.3	82.0	**88.1**	**90.9**	**77.0**	99.8	**37.4**	**30.2**

	Planning				Pose refinement		
	PointGoal		ObjGoal		Conv.	Rot.	Trans.
Policy	Succ.	SPL	Succ.	SPL	rate	error (°)	error (m)
Frontier	22.4	21.4	9.6	9.1	7.2	0.383	0.00955
E2E (cov.)	30.0	29.3	8.9	8.3	15.4	0.319	0.00775
E2E (cur.)	29.8	29.2	8.5	8.0	12.5	0.325	0.00799
E2E (nov.)	32.3	31.9	11.4	10.8	19.4	0.315	0.00774
E2E (rec.)	32.8	32.6	10.5	10.0	19.3	0.292	0.00734
Ours (cov.)	**39.5**	**39.0**	14.8	14.3	20.2	**0.283**	**0.00734**
Ours (sem.)	37.7	37.4	**16.0**	**15.4**	**23.0**	0.319	0.00784
Ours (obs.)	38.2	37.8	**15.8**	**15.3**	22.5	0.305	0.00765
Ours (view.)	39.0	38.6	**15.9**	**15.3**	21.1	0.316	0.00769