COMPASS: Learning Global Spatial Context for Long-Range Robot Navigation

Jose Lavariega-Gomez[1,2], Nicola Irmiger[1,2], Erica Tevere[1], Patrick Spieler[1]

***Abstract*— Autonomous long-range navigation in unstructured environments is challenging due to limited sensing range, partial observability, and the difficulty of estimating the long-term utility of unexplored regions. We present COMPASS, a visual-geometric navigation framework that learns a goal-conditioned frontier utility policy over graph-structured spatial memory and semantic visual observations. COMPASS combines a Graph Convolutional Network (GCN) operating over a persistent navigation graph with DINOv3 visual embeddings extracted from map frontier observations. The graph captures global topology and traversability structure, while visual embeddings provide non-local cues regarding paths, openings, and terrain continuity. A Deep Q-Learning policy estimates frontier utility under delayed rewards, enabling behaviors such as backtracking and dead-end recovery. To train the policy efficiently, we introduce a lightweight graph-based rollout simulator that expands limited field data into millions of navigation experiences. Experimental results in unseen outdoor environments demonstrate improved long-range navigation performance compared to heuristic frontier-selection baselines.**

I. INTRODUCTION

Autonomous long-range navigation in unknown outdoor environments remains challenging due to limited sensing range, partial observability, and the delayed consequences of exploration decisions. Without a prior map, a robot must repeatedly select frontiers at the boundary of explored space while reasoning about dead ends, backtracking, and future reachability. Traditional frontier-based exploration methods typically rely on handcrafted utility functions or local geometric heuristics that become myopic in large outdoor environments with sparse semantic structure and frequent occlusions. Purely geometric methods often lack semantic foresight, while purely visual approaches may fail to preserve globally consistent spatial memory over long horizons [1].

Prior works [2], [3], [4] have combined graph memory with visual reasoning in semantically rich indoor environments, but often rely on handcrafted frontier priors, task-specific heuristics, or extensive simulator interaction. In contrast, we investigate whether frontier utility can be learned directly from navigation experience by jointly reasoning over graph topology and semantic visual observations.

[1]Jet Propulsion Laboratory (JPL), California Institute of Technology (Caltech), Pasadena, CA 91011, USA. Corresponding author: Patrick Spieler (patrick.spieler@jpl.nasa.gov).

[2]Swiss Federal Institute of Technology (ETH Zürich).

The research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration (80NM0018D0004). The work was supported in part by the Luxembourg National Research Fund (Ref. 18990533) and the Swiss National Science Foundation as part of projects No. 200021E_229503 and No. 227617. Jose Lavariega was supported by the Heyning Roelli Foundation and the Swiss-European Mobility Program during this work.

Our work, COMPASS is a long-range visual-geometric navigation framework that learns a goal-conditioned frontier utility policy over graph-structured spatial memory. COMPASS combines a persistent LiDAR-based navigation graph with semantic embeddings extracted from frontier imagery using the pretrained DINOv3 foundation model. A Graph Convolutional Network (GCN) encodes global topology and traversability context, while a Deep Q-Learning policy estimates frontier utility under delayed rewards, enabling recovery behaviors such as dead-end backtracking.

II. METHODOLOGY

COMPASS maintains a persistent geometric navigation graph constructed from LiDAR observations that acts as a long-term spatial memory throughout exploration. Nodes encode traversable robot poses, while edges represent safe transitions between explored regions. RGB images and structure estimates are additionally used to build multi-view frontier observations from explored space, reducing occlusion ambiguity.

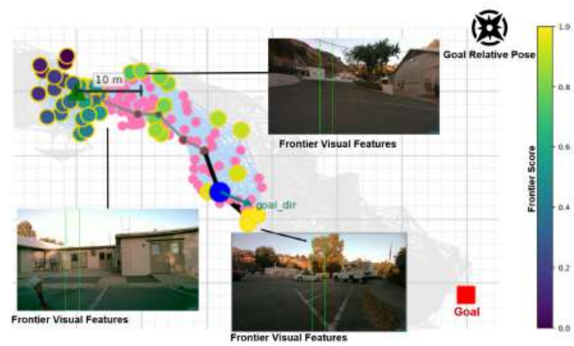

Fig. 1. COMPASS Observation Space: A navigation graph with images of the frontiers as seen from multiple viewpoints within the explored space.

We jointly reason over graph topology and semantic frontier observations using a Graph Convolutional Network (GCN) over the navigation graph and DINOv3 [5] semantic visual embeddings extracted from frontier imagery. The GCN encodes graph topology, traversability structure, node connectivity, and local geometry to produce a goal-conditioned latent representation estimating long-range cost-to-go, even when the goal is not yet connected to explored space. Frontier image features are obtained by projecting graph frontier positions onto the corresponding feature image, allowing the embeddings to capture semantic cues associated with paths, openings, terrain continuity, and visually traversable regions.

Both modalities are fused by spatially matching graph frontiers with their corresponding DINO embeddings (Fig. 1). Additional geometric features, including goal direction, distance-to-goal, and outgoing edge traversability, are incorporated into a Deep Q-Learning framework that estimates frontier utility under delayed rewards.

To train efficiently from limited field data, we construct a lightweight graph-based rollout simulator using traversals collected across hiking trails, urban areas, and moderately vegetated environments. Several hours of real-world data are expanded into millions of randomized navigation experiences. We use an adaptive curriculum with progressively increasing start-goal separation, to learn long-horizon selection and recovery.

COMPASS builds upon prior graph-based exploration systems [6] that score frontiers using manually engineered utility functions. In contrast, our approach learns frontier utility directly from navigation experience through reinforcement learning over graph-structured spatial memory and semantic visual observations. Rather than relying on handcrafted frontier priors, the policy learns navigation-relevant semantic and geometric patterns associated with successful long-range exploration behaviors across diverse outdoor environments.

III. RESULTS

We evaluate COMPASS in simulation and real-world outdoor deployments across urban and natural environments. Simulation episodes use randomized start-goal pairs from unseen environments with trajectories up to 300m. Performance differences between learned and heuristic frontier-selection methods increase with navigation distance and dead-end recovery requirements.

We further deploy the policy on a quadruped platform equipped with LiDAR and RGB cameras across urban streets, open fields, hills, and hiking trails.

$$S_{rate} = \frac{1}{E}\sum_{i=1}^{E} S_i \quad (1) \qquad SPL = \frac{1}{E}\sum_{i=1}^{E} \frac{S_i l_i}{\max(l_i, p_i)} \quad (2)$$

$$Cost_{norm} = \frac{1}{|S|}\sum_{i \in S} \frac{p_i}{l_i} \quad (3)$$

We report success rate (Eq. 1), success weighted by path length (Eq. 2), and normalized path cost among successful episodes (Eq. 3), where S_i indicates goal completion, E is the number of evaluation episodes, and l_i, p_i denote optimal and traversed path cost respectively.

COMPASS improves success rate and path efficiency relative to heuristic frontier-selection baselines [1], particularly in environments with dead ends and limited geometric visibility. In graph-based simulation over 150 long-range trajectories (150-300m), COMPASS achieves up to a 20% success-rate improvement over variants without graph or semantic visual features (Fig. 2).

Real-world experiments demonstrate that persistent graph memory and periodic replanning allow adaptation to transient obstacles and changing environments (Fig. 3). Current limitations include the computational cost of frontier-image encoding, storage and the growing frontier action space during extended exploration.

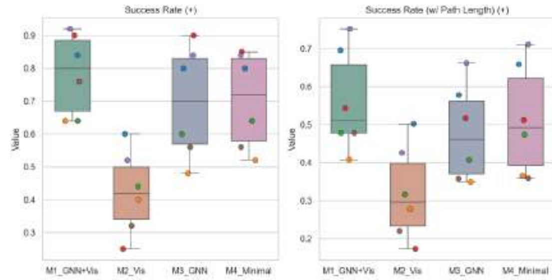

Fig. 2. Success rate and SPL evaluation across outdoor navigation datasets comparing COMPASS against graph-only, visual-only, and heuristic frontier-selection baselines. Episodes terminate after 50 frontier expansions. Colored dots correspond to the same set of evaluation trajectories captured from different datasets.

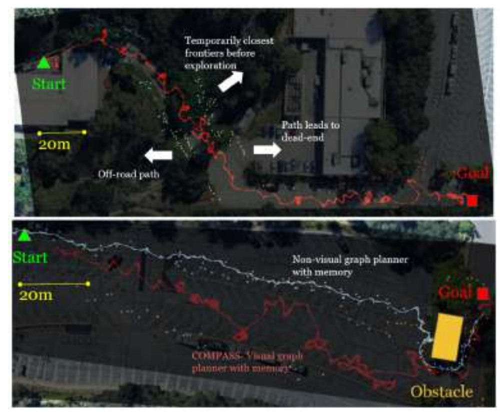

Fig. 3. (Top) Hardware trajectory successfully reaching a goal 300m away, in the presence of a dynamic obstacle, dead ends, and multiple alternate paths. (Bottom) comparison between our visually-augmented method against a planner without visual augmentation (white). The visual planner is able to reason about an opening in the image frame and makes a more direct path towards the opening.

IV. CONCLUSION

We presented COMPASS, a long-range visual-geometric navigation framework that learns frontier utility over graph-structured spatial memory and semantic visual observations. Experimental results demonstrate improved long-range navigation performance compared to heuristic frontier-selection methods, highlighting the potential of learned frontier utility estimation for autonomous navigation in outdoor environments.

REFERENCES

[1] M. Schmittle, *et al.*, “Long range navigator (lrn) : Extending robot planning horizons beyond metric maps,” *arXiv*, 2025.

[2] H. Li, *et al.*, “Memonav: Working memory model for visual navigation,” 2024. [Online]. Available: https://arxiv.org/abs/2402.19161

[3] L. Mezghan, *et al.*, “Memory-augmented reinforcement learning for image-goal navigation,” in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2022, pp. 3316–3323.

[4] O. Kwon, *et al.*, “Visual graph memory with unsupervised representation for visual navigation,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2021, pp. 15 890–15 899.

[5] O. Siméoni, *et al.*, “Dinov3,” 2025. [Online]. Available: https://arxiv.org/abs/2508.10104

[6] H. Shah, *et al.*, “Wildos: Open-vocabulary object search in the wild,” 2026. [Online]. Available: https://arxiv.org/abs/2602.19308