

# An Efficient Beam Search Algorithm for Active Perception in Mobile Robotics

Kaixian Qu, Han Wang, Victor Klemm, Cesar Cadena, Marco Hutter

**Abstract**—Active perception is a fundamental problem in autonomous robotics, in which the robot must decide where to move and what to sense in order to obtain the most informative observations for accomplishing its mission. Existing approaches often solve a computationally expensive traveling salesman problem over heuristically selected nodes, or adopt an overly constrained shortest path tree formulation. We introduce node-wise beam search (NBS), an efficient algorithm that maintains the top candidate paths per node to effectively explore the solution space. To balance exploration and exploitation, we integrate frontier information into a novel expected gain metric. Furthermore, we propose the rapidly-exploring random annulus graph (RRAG), an incremental graph construction method that preserves full orientation sampling and ensures connectivity in cluttered environments via a fallback local planner. Evaluations across three active perception tasks demonstrate that NBS combined with RRAG outperforms state-of-the-art baselines by at least 20% in one or more tasks.

## I. INTRODUCTION

Autonomous robots are increasingly deployed for active perception tasks requiring intelligent information gathering [1], [2]. Sampling-based tree planners suffer from heavy rewiring, pruning, and re-expanding overhead, as well as the constraint of storing only one path per node [3], [4]. Graph-based methods offer multi-query capabilities but often rely on restrictive shortest path trees (SPT) [5] or decouple routing using computationally expensive traveling salesman problem (TSP) solvers [6], [7]. To address these limitations, we introduce node-wise beam search (NBS), which maintains multiple top candidate paths per node rather than per search depth. Integrated with the proposed expected gain metric and the rapidly-exploring random annulus graph (RRAG), NBS outperforms other state-of-the-art algorithms by a significant margin.

## II. METHODOLOGY

### A. Path Selection and Expected Gain

To explicitly account for exploration, we introduce the *expected gain* metric, incorporating frontier information [8]. For a path  $p$  terminating at a frontier node  $\mathcal{P}_{\mathcal{F}}$ , we approximate future gain by assuming the robot continues into the unmapped region at its current gain-to-cost ratio over the remaining budget  $C$ :

$$g_e(p) = \begin{cases} r(p) \cdot C, & \text{if } p \in \mathcal{P}_{\mathcal{F}}, \\ g(p), & \text{otherwise.} \end{cases} \quad (1)$$

The authors are with the Robotic Systems Lab, ETH Zurich. [kaiqu@ethz.ch](mailto:kaiqu@ethz.ch)

This extended abstract summarizes work that has been conditionally accepted to the International Journal of Robotics Research (IJRR).

Here,  $c(p)$  denotes the cost of a path,  $g(p)$  the gain of a path, and  $r(p) = g(p)/c(p)$  represents the path ratio. The robot continuously replans at every visited node to adapt to newly acquired information.

### B. Node-wise Beam Search (NBS)

Standard depth-wise beam search (DBS) is prone to local optima if the entire beam converges on high immediate gains. NBS (Algorithm 1) overcomes this by maintaining up to  $B$  candidate paths for each visited node. This guarantees that at least one promising path remains available across the graph, ensuring robust performance even with a minimal beam width of  $B = 1$ . The heap structure ( $open[v]$ ) stores up to the top  $B$  paths, prioritizing those with a better gain-to-cost ratio, followed by path gain, and finally by lower cost. NBS exhibits a total time complexity of  $\mathcal{O}(|\mathcal{E}|DB(D + \log B))$  and a space complexity of  $\mathcal{O}(|\mathcal{V}|DB)$ .

---

#### Algorithm 1: NBS (simplified)

---

**Input:** Graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ , start node  $v_s$   
**Param:** Beam width  $B$ , search depth  $D$ , quality  $g_e$

```

1  $p^* \leftarrow (v_s)$ ,  $open[v_s] \leftarrow \{(v_s)\}$ 
2 for  $d = 1$  to  $D$  do
3    $heap \leftarrow \emptyset$ 
4   foreach  $v \in \mathcal{V}$  do
5     foreach  $p \in open[v]$  do
6       foreach  $(v, v') \in \mathcal{E}$  do
7         if  $edge(v, v')$  not in  $p$  then
8            $p' \leftarrow p + (v, v')$ 
9           if  $c(p') \leq C$  then
10             if  $g_e(p') > g_e(p^*)$  then
11                $p^* \leftarrow p'$ 
12             Insert  $p'$  into  $heap[v']$ ,
              retaining at most  $B$  paths
13    $open \leftarrow heap$ 
14 return  $p^*$ 

```

---

### C. Graph Construction via RRAG

The robot incrementally constructs the rapidly-exploring random annulus graph (RRAG). Collision-free configurations are accepted if their distance to all existing nodes exceeds  $l_{\min}$ , with edges formed if the distance falls within  $[l_{\min}, l_{\max}]$ . To prevent the graph from stalling in cluttered spaces where straight-line interpolation fails, RRAG

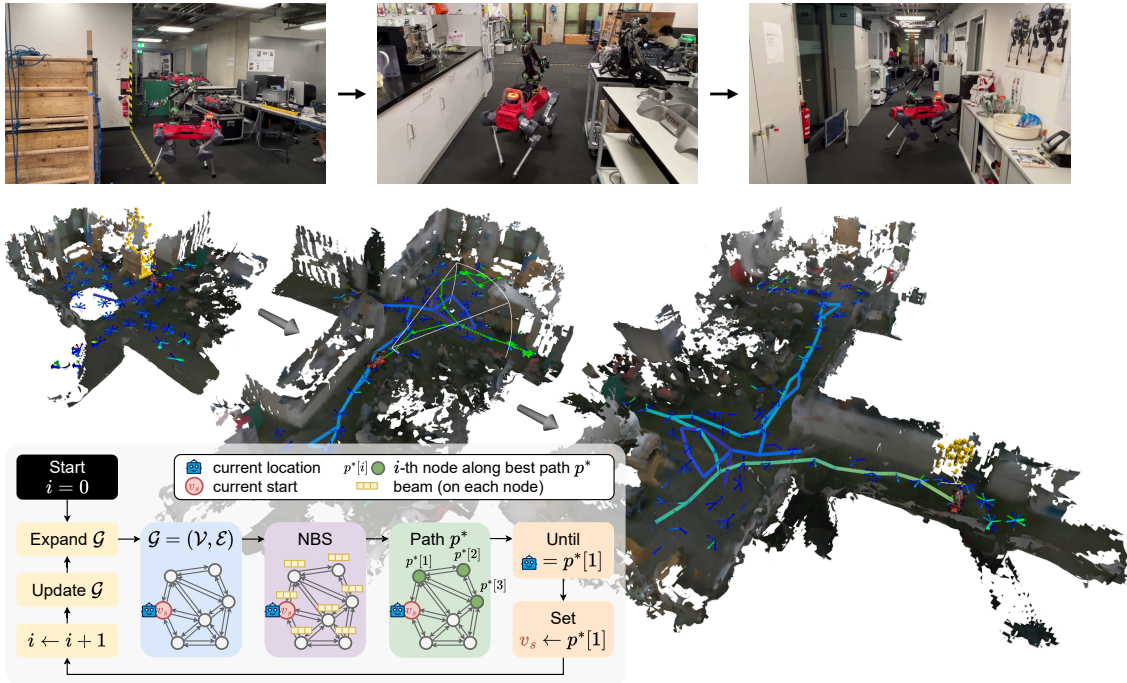


Fig. 1. **Surface reconstruction in hardware experiments.** The planner identifies the optimal path using NBS and updates the graph before computing the next-best path. The trajectory is visualized with a color gradient; yellow voxels represent visible surface frontiers.




Fig. 2. **Comparison of normalized gain across methods and tasks.** NBS combined with RRAM achieves the highest performance across all active perception tasks.

integrates a fallback local sampling-based planner (FLS) utilizing RRT\* [9]. Furthermore, RRAM preserves multiple discrete orientations as distinct nodes at each position, preventing inefficient re-orientation maneuvers.

### III. EXPERIMENTS AND RESULTS

We evaluated our framework in Habitat [10] across eight scenarios from the HSSD dataset [11]. We benchmarked NBS against DBS, SPT/DEP [5], TSP [7], FUEL [6], RRT\* [4], and RRT [3]. Tasks included: (1) **point collection**, (2) **surface reconstruction** [12], and (3) **volumetric exploration**.

#### A. Simulation Performance

As shown in Fig. 2, NBS combined with RRAM achieved the highest performance across all tasks. In the constrained point collection task, NBS achieved a 0.50 normalized gain, outperforming the second-best algorithm by 16%. While DBS performed well in surface reconstruction, it was highly

susceptible to local optima in sparser tasks. TSP-based algorithms suffered from excessive computational demands, often failing to exploit the cost budget within the allotted time limit.

#### B. Hardware Validation

We deployed our pipeline on an ANYmal quadruped [13] equipped with a ZED X Mini camera (Fig. 1). Computations, including Voxblox [14] mapping and NBS planning, ran fully onboard. We further tested our planner across diverse environments, including an outdoor space, an engine museum, a cafeteria, and various laboratory settings.

### IV. CONCLUSION

NBS combined with RRAM forms an efficient framework for active perception. By maintaining multiple paths per node, utilizing an expected gain criterion, and ensuring robust graph connectivity, the system elegantly balances exploration and exploitation, consistently outperforming state-of-the-art baselines.

## REFERENCES

- [1] N. Cao, K. H. Low, and J. M. Dolan, "Multi-robot informative path planning for active sensing of environmental phenomena: A tale of two algorithms," *arXiv preprint arXiv:1302.0723*, 2013.
- [2] M. Popović, T. Vidal-Calleja, G. Hitz, J. J. Chung, I. Sa, R. Siegwart, and J. Nieto, "An informative path planning framework for uav-based terrain monitoring," *Autonomous Robots*, vol. 44, no. 6, pp. 889–911, 2020.
- [3] A. Bircher, M. Kamel, K. Alexis, H. Oleynikova, and R. Siegwart, "Receding horizon" next-best-view" planner for 3d exploration," in *2016 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2016, pp. 1462–1468.
- [4] L. Schmid, M. Pantic, R. Khanna, L. Ott, R. Siegwart, and J. Nieto, "An efficient sampling-based method for online informative path planning in unknown environments," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 1500–1507, 2020.
- [5] Z. Xu, D. Deng, and K. Shimada, "Autonomous uav exploration of dynamic environments via incremental sampling and probabilistic roadmap," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 2729–2736, 2021.
- [6] B. Zhou, Y. Zhang, X. Chen, and S. Shen, "Fuel: Fast uav exploration using incremental frontier structure and hierarchical planning," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 779–786, 2021.
- [7] S. Arora and S. Scherer, "Randomized algorithm for informative path planning with budget constraints," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2017, pp. 4997–5004.
- [8] B. Yamauchi, "A frontier-based approach for autonomous exploration," in *Proceedings 1997 IEEE International Symposium on Computational Intelligence in Robotics and Automation CIRA'97. Towards New Computational Principles for Robotics and Automation'*. IEEE, 1997, pp. 146–151.
- [9] S. Karaman and E. Frazzoli, "Sampling-based algorithms for optimal motion planning," *The international journal of robotics research*, vol. 30, no. 7, pp. 846–894, 2011.
- [10] M. Savva, A. Kadian, O. Maksymets, Y. Zhao, E. Wijmans, B. Jain, J. Straub, J. Liu, V. Koltun, J. Malik, D. Parikh, and D. Batra, "Habitat: A platform for embodied ai research," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2019, pp. 9339–9347.
- [11] M. Khanna, Y. Mao, H. Jiang, S. Hareesh, B. Shacklett, D. Batra, A. Clegg, E. Undersander, A. X. Chang, and M. Savva, "Habitat synthetic scenes dataset (hssd-200): An analysis of 3d scene scale and realism tradeoffs for objectgoal navigation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 16 384–16 393.
- [12] L. Yoder and S. Scherer, "Autonomous exploration for infrastructure modeling with a micro aerial vehicle," in *Field and Service Robotics: Results of the 10th International Conference*. Springer, 2016, pp. 427–440.
- [13] M. Hutter, C. Gehring, D. Jud, A. Lauber, C. D. Bellicoso, V. Tsounis, J. Hwangbo, K. Bodie, P. Fankhauser, M. Bloesch, R. Diethelm, S. Bachmann, A. Melzer, and M. Hoepflinger, "Anymal-a highly mobile and dynamic quadrupedal robot," in *2016 IEEE/RSJ international conference on intelligent robots and systems (IROS)*. IEEE, 2016, pp. 38–44.
- [14] H. Oleynikova, Z. Taylor, M. Fehr, R. Siegwart, and J. Nieto, "Voxblox: Incremental 3d euclidean signed distance fields for on-board mav planning," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2017, pp. 1366–1373.