Human-Aware Goal-Oriented Exploration Using Frontier-Biased RRT and Vision-Conditioned Trajectory Selection

Muhammad Imad* and Soo Young Shin†*
IT Convergence Engineering, Kumoh National Institute of Technology,
Daehak-ro 61, Gumi-si, (39177) Gyeongbuk, South Korea
Email: imadsafi08@kumoh.ac.kr, wdragon@kumoh.ac.kr

***Abstract*—This paper presents a human-aware exploration framework that integrates a modified frontier-based rapidly-exploring random tree planner with a vision-conditioned local navigation policy. Frontier exploration is reformulated using forward-biased sampling combined with heading alignment constraints and a clearance-aware goal backoff mechanism that retracts navigation targets from unknown boundaries to improve local feasibility. Motion execution is performed using a Neural Value-Driven Policy, which samples multiple candidate trajectories conditioned on visual observations and selects control actions using a learned critic function. To support semantic mission objectives, a robust identity-gating mechanism enforces temporal persistence and spatial consistency before confirming human detections, thereby preventing duplicate counting under tracker instability. Exploration continues until a predefined semantic objective is satisfied, after which the system transitions to a return-home behavior.**

***Index Terms*—Autonomous Exploration, Frontier-Based Planning, Rapidly-Exploring Random Trees, Vision-Based Navigation, Human-Aware Robotics**

I. Introduction

Autonomous exploration has traditionally been formulated as a geometric problem, where the objective is to maximize map coverage or information gain in unknown environments. Frontier-based exploration methods explicitly target the boundary between known and unknown space and have demonstrated strong performance in practice [2]. Sampling-based planners such as rapidly-exploring random trees provide an efficient mechanism for exploring high-dimensional spaces under differential constraints [1].

In many real-world applications including search and rescue and service robotics, exploration must be guided by semantic objectives rather than geometric coverage alone. In such scenarios, the robot must reliably detect and reason about humans. Naively combining exploration with human detection often results in repeated counting caused by identity flickering, short-lived false positives, and tracker instability [4]. In addition, navigation goals selected directly on frontier boundaries frequently lead to execution failure due to insufficient clearance near unknown regions.

This work addresses these challenges by embedding feasibility constraints and semantic reasoning directly into the exploration loop. The proposed system integrates a frontier-biased RRT planner with a learned vision-conditioned navigation policy and a robust identity confirmation mechanism, enabling reliable human-aware exploration with deterministic mission termination.

II. System Overview

The proposed framework is structured hierarchically to separate long-horizon decision making from short-horizon motion execution, while maintaining tight coupling through shared objectives. An overview of the system architecture and information flow is illustrated in Fig. 1.

Let the robot state at time t be denoted as $x_t = (p_t, \theta_t)$, where $p_t \in \mathbb{R}^2$ represents the planar position and θ_t the heading. The environment is represented as a two-dimensional occupancy grid $\mathcal{M}$, a standard abstraction in probabilistic robotics [3]. A frontier set $\mathcal{F}$ is extracted as the boundary between known free space and unknown regions.

Global exploration operates by growing a rapidly-exploring random tree $\mathcal{T}$ rooted at the current robot state. Sampling is biased toward the robot forward direction to encourage motion continuity:

$$x_s \sim \begin{cases} \mathcal{C}(x_t, \theta_t), & \text{with probability } p_f, \\ \mathcal{U}(\mathcal{M}), & \text{otherwise,} \end{cases} \tag{1}$$

where $\mathcal{C}(\cdot)$ denotes a bounded forward sampling cone and $\mathcal{U}(\cdot)$ uniform sampling over the map domain.

Candidate nodes are accepted only if they are collision free and lie within a bounded distance of a frontier cell, following frontier-driven exploration principles [2]. To improve feasibility near unknown regions, selected frontier goals are retracted by a standoff distance d_s along the approach direction and must satisfy a free-disk constraint of radius r_c:

$$\forall c \in \mathcal{D}(r_c), \quad \mathcal{M}(c) = 0. \tag{2}$$

Once a global goal g_k is selected, local motion execution is delegated to a Neural Value-Driven Policy. Given a visual observation o_t, the policy samples a set of K candidate trajectories:

$$\{\tau_t^{(i)}\}_{i=1}^{K} \sim \pi_\theta(\cdot \mid o_t), \tag{3}$$

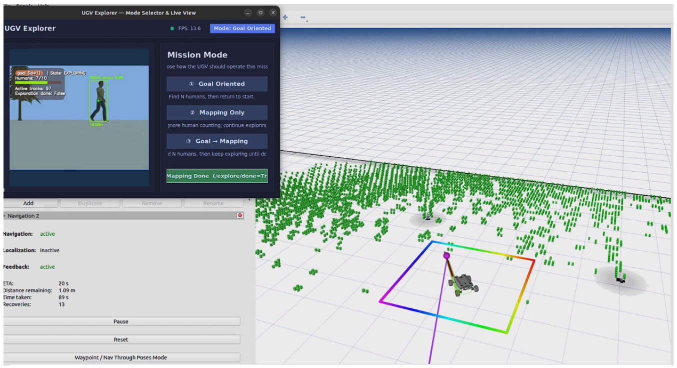

Fig. 1. Overview of the proposed human-aware exploration framework integrating frontier-biased RRT planning, vision-conditioned local trajectory selection, and robust human identity gating within a unified mission loop.

where each trajectory is a finite sequence of velocity commands. A learned critic function Q_ϕ evaluates each candidate, and the selected action is obtained as

$$\tau_t^* = \arg\max_{\tau_t^{(i)}} Q_\phi(o_t, \tau_t^{(i)}). \tag{4}$$

III. Human Identity Gating

Human detections are accumulated over time and associated with tracked identities. Each detected human is represented by an observation state

$$\mathcal{O}_i = \{t_{first}, t_{last}, k_i, (x_i, y_i)\}, \tag{5}$$

where k_i denotes the number of consecutive detections.

An identity is confirmed only if the temporal stability condition

$$k_i \geq k_{\min} \ \wedge \ (t_{last} - t_{first}) \geq T_{\min} \tag{6}$$

is satisfied. This design mitigates transient detections and tracker noise [4].

To prevent identity switching and duplicate counting, a candidate identity is rejected if

$$\exists h \in \mathcal{H}_{recent} : \|(x_i, y_i) - (x_h, y_h)\| \leq r_{flip} \wedge (t_i - t_h) \leq T_{flip}. \tag{7}$$

Once confirmed, identities are permanently blocked from re-entry, ensuring consistent human counting throughout the mission.

IV. Mission-Level Algorithm

Algorithm 1 Human-Aware Exploration Mission

Initialize occupancy map $\mathcal{M}$ and confirmed human set $\mathcal{H} \leftarrow \emptyset$
Set target human count N_H
while $|\mathcal{H}| < N_H$ **do**
  Extract frontier set $\mathcal{F}$ from $\mathcal{M}$
  Grow frontier-biased RRT $\mathcal{T}$ from current state [1]
  Select frontier goal g_k satisfying heading and clearance constraints
  while goal g_k not reached **do**
    Acquire visual observation o_t
    Sample candidate trajectories $\{\tau_t^{(i)}\}_{i=1}^K$
    Select $\tau_t^* = \arg\max Q_\phi(o_t, \tau_t^{(i)})$
    Execute control command from τ_t^*
    Update map $\mathcal{M}$
    Update human observation states $\{\mathcal{O}_i\}$
    if identity i satisfies stability constraints **then**
      $\mathcal{H} \leftarrow \mathcal{H} \cup \{i\}$
    end if
  end while
end while
Navigate to home location g_{home}
Terminate mission

V. Conclusion

This paper presented a human-aware exploration framework that integrates frontier-biased RRT planning, vision-conditioned trajectory selection, and robust semantic reasoning. By modifying frontier exploration with clearance-aware goal retraction and coupling it with a learned value-driven navigation policy, the system achieves reliable exploration and human counting in unknown environments. The proposed approach provides a principled foundation for future extensions to multi-robot exploration and additional semantic objectives.

Acknowledgment

This work was supported by the National Research Foundation of Korea(NRF) grant funded by the Korea government(MSIT) (RS-2025-00553810, 50%) This research was supported by the MSIT(Ministry of Science and ICT), Korea, under the ITRC(Information Technology Research Center) support program(IITP-2025-RS-2024-00437190, 50%) supervised by the IITP(Institute for Information & Communications Technology Planning & Evaluation)

References

[1] S. M. LaValle, "Rapidly-exploring random trees: A new tool for path planning," IEEE Technical Report, 1998.
[2] B. Yamauchi, "A frontier-based approach for autonomous exploration," in *Proceedings of the IEEE International Symposium on Computational Intelligence in Robotics and Automation*, 1997.
[3] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics*. MIT Press, 2005.
[4] A. Bewley *et al.*, "Simple online and realtime tracking," in *Proceedings of the IEEE International Conference on Image Processing*, 2016.
[5] M. Quigley *et al.*, "ROS: an open-source robot operating system," in *Proceedings of the IEEE International Conference on Robotics and Automation Workshop*, 2009.