

Autonomous Robot for Disaster Mapping and Victim Localization

Donatello Goes On A Rescue Mission

Michael Potter
Cognitive Systems Lab
Northeastern University
Boston, USA
potter.mi@northeastern.edu

Rahil Bhowal
Masters in Robotics
Northeastern University
Boston, USA
bhowal.r@northeastern.edu

Richard Zhao
Masters in Computer Science
Northeastern University
Boston, USA
zhao.richa@northeastern.edu

Anuj Patel
Masters in Robotics
Northeastern University
Boston, USA
patel.anuj2@northeastern.edu

Jingming Cheng
Masters in Robotics
Northeastern University
Boston, USA
cheng.jingm@northeastern.edu

Abstract—In response to the critical need for effective reconnaissance in disaster scenarios, this research article presents the design and implementation of a complete autonomous robot system using the Turtlebot3 Burger (TB3) with Robot Operating System (ROS) Noetic. Upon deployment in closed, initially unknown environments, the system aims to generate a comprehensive map and identify any present 'victims' using AprilTags as stand-ins. We discuss our solution for search and rescue missions, while additionally exploring more advanced algorithms to improve search and rescue functionalities. We introduce a Cubature Kalman Filter (CKF) to help reduce the Mean Squared Error (MSE) [m] for AprilTag localization and an information-theoretic exploration algorithm to expedite exploration in unknown environments. Just like turtles, our system takes it slow and steady, but when it's time to save the day, it moves at ninja-like speed! Despite Donatello's shell, he's no slowpoke - he zips through obstacles with the agility of a teenage mutant ninja turtle. So, hang on tight to your shells and get ready for a whirlwind of reconnaissance!

Full pipeline code <https://github.com/rzhao5659/MRProject/tree/main>
Exploration code <https://github.com/rzhao5659/MRProject/tree/rz/exploration>

Index Terms—reconnaissance, ROS, Donatello, Turtlebot3, April-Tag

I. PROBLEM STATEMENT

In the wake of disasters, rapid and efficient reconnaissance is critical for successful rescue operations. Traditionally, human responders perform this task, but they often face considerable risks and limitations in accessibility. The development of autonomous robotic systems presents a promising alternative, capable of navigating and mapping disaster environments quickly and safely. However, existing autonomous systems

primarily focus on mapping and fail to effectively integrate victim identification within their operational protocols.

This project we addresses the need for an autonomous robotic system that can both map unknown, closed environments and identify "victims(AprilTags)" in these settings. Our goal is to design and implement a robust system using the Robot Operating System (ROS) Noetic, which can:

- Generate a complete occupancy grid map of an unknown environment, providing essential data for navigation and strategic planning.
- Locate and accurately estimate the poses of "victims," represented by AprilTags, to simulate the identification and location of individuals in need of rescue.

The system must operate entirely autonomously, adapting to the dynamic and unpredictable nature of disaster environments. The key challenges include integrating advanced mapping techniques with efficient victim detection algorithms, ensuring the system's reliability in diverse conditions, and managing computational constraints potentially imposed by the hardware used.

II. SETUP

The PC setup, Single Board Computer setup, and OpenCR setup were completed by adhering to the Turtlebot3 Burger (TB3) setup instructions provided on the Robotis website [1]. Subsequently, we progressed to more advanced software and hardware configurations after the initial setup.

A. TurtleBot3 Hardware Setup

The TB3 is a widely used, compact, and customizable robotic platform suitable for a variety of robotic applications, par-

ticularly in educational and research environments. For our project, the TB3 platform was selected due to its modularity, ease of use, and compatibility with ROS. Below, we detail the hardware setup and configuration tailored for our autonomous reconnaissance system.

1) *Robot Model*: We employed the TB3 [model](#) for our project. This model is particularly advantageous for indoor environments due to its compact size and maneuverability. The key components of the TB3 include:

- Dynamixel X-series motors: These servos provide precise and reliable motion control.
- [360-degree LiDAR sensor\(LDS-02\)](#): Essential for accurate mapping and obstacle avoidance.
- [Raspberry Pi](#): Serves as the main processing unit, capable of handling both ROS operations and additional computation for sensor data processing.

2) *Sensor Configuration*: To enhance the TB3's capabilities for detailed environmental mapping and victim identification, we configured the following sensor setup:

- [LiDAR](#): Mounted at the top of the robot, providing a 360-degree field of view to detect obstacles and assist in generating a high-resolution occupancy grid map.
- [Camera](#): Raspberry Pi Camera was attached to capture visual data, crucial for AprilTag detection. This camera is positioned to have an unobstructed view of the environment ahead of the robot.

3) *Additional Hardware*: Given the computational demands of our project, particularly for processing occupancy grids and detecting AprilTags, we supplemented the onboard Raspberry Pi with the following:

- External Battery Pack: To ensure extended operational periods without the need for recharging.
- Wi-Fi Adapter: To enhance the communication capabilities with the ROS network and enable off-board processing if required.

4) *Calibration*: Before deployment, the TB3 underwent a series of calibration procedures to ensure accuracy in its mapping and navigation capabilities:

- LiDAR Calibration: To verify the accuracy of distance measurements and angular resolution.
- Camera Calibration: Essential for precise AprilTag detection, involving adjusting the focus and verifying the camera's field of view aligns with the expected operational environment. We followed the ROS tutorials for monocular camera calibration [2].

B. TurtleBot3 Software Setup

Both TurtleBot and remote PC run on ROS Noetic. For communication between these two, they need to be connected to the same network, and the remote PC need to run **roscore** to act as the ROS master.

On the TurtleBot, it's necessary to install the packages `TurtleBot3_description`, `TurtleBot3_bringup`, `raspicam_node`, and `apriltag_ros`. Our solution launches these files on the TurtleBot for bringup and AprilTags detection:

- `TurtleBot3_robot.launch` from `TurtleBot3_bringup`.
- `camerav2_640_624_15fps.launch` from `raspicam_node`.
- `continuous_detection.launch` from `apriltag_ros`.

The `TurtleBot3_robot.launch` starts up the RPLidar and the OpenCR board. It publishes all sensor topics and subscribes to `/cmd_vel` for controlling the wheel motors. It also provides wheel odometry as a transform from frame `odom` to frame `base_footprint`, and as a message in `/odom`.

The `camerav2_640_624_15fps.launch` starts up the Raspberry Pi Camera Module in 640x624 resolution with 15 Frames Per Second (FPS) framerate. This publishes the raw image in `/image_raw`, and the camera calibration matrix in `/camera_info`.

The `continuous_detection.launch` uses the raw camera output to detect and estimate AprilTag pose. It provides a transform from camera frame `raspicam` to the tag frame `Tag<id>`. This information is also published in `/tag_detections`.

On the remote PC, it's necessary to install the external packages `TurtleBot3_description`, `TurtleBot3_navigation`, `TurtleBot3_simulations`, `TurtleBot3_slam`, and `m-explore`. Our solution launches these files on the remote PC to launch the autonomy stack:

- `TurtleBot3_slam.launch` from `TurtleBot3_slam`.
- `move_base.launch` from `TurtleBot3_navigation`.
- `explore_lite.launch` from `m-explore`.
- `ckf3D` node.
- `search and rescue` node.

The `TurtleBot3_slam.launch` launches gmapping to perform Simultaneous Localization and Mapping (SLAM) using LIDAR scans, and a robot state publisher node to publish all sensor transforms relative to the robot frame `base_link`, which are specified in the URDF file. The gmapping node provides an occupancy grid map in `/map` as well as a transform from frame `map` to frame `odom`. This completes the transform tree, and allows us to obtain the robot's current pose relative to the map frame.

The `move_base.launch` launches move base node, which performs planning and control. It subscribes to an action topic `move_base/goal` that specifies the goal, plans a path using the default global and local planners, and publishes back the command velocity to `/cmd_vel` topic, which the TB3 listens to. The move base utilizes the occupancy grid map `/map` obtained from SLAM and inflates the obstacles seen to produce a costmap, where the planning is performed.




Fig. 1. Gazebo simulation arenas. (a) is $15.0\text{m} \times 10.0\text{m}$ (b) is $3.0\text{m} \times 3.0\text{m}$ (c) is $34.5\text{m} \times 21.75\text{m}$.

The *explore_lite.launch* runs a greedy frontier-based algorithm to compute the next goal for exploration. It receives the occupancy grid map either from gmapping or the costmap from *move_base*, and publishes the exploration goal in *move_base/goal*.

These three nodes complete the essential autonomy stack and enables our robot to explore and map an environment autonomously, thus completing the first task of the problem. The remaining two nodes are used for the second task of the problem, which is identification and pose estimation of the "victims" (represented as AprilTags).

The *ckf3D* node launches a Cubature Kalman Filter (CKF) [3] that continuously receives the last pose estimates from the AprilTagdetector in */tag_detections* and outputs a filtered pose estimate for each identified tag as new transforms. This improves the accuracy of pose estimate for the second task.

The *search_and_rescue* node is the main algorithm that runs after the exploration node finishes its task, indicated through a */exploration_stop_flag* topic. This computes a series of goals that covers uniformly the explored map and order them in a queue such that the robot would move in a snake-like (zig zag) search pattern. These will then be sent to move base through the action topic *move_base/goal*. In addition, this node directly commands the TurtleBot to spin after reaching each goal. This node will allow the robot to identify any missing AprilTag during the exploration, and improve the pose estimation of AprilTags by increasing the number of pose measurements for the *ckf3D* node. Once the queue is empty and all the AprilTags are processed, the *search_and_rescue* node publishes the */finish_search* topic, to which the *ckf3D* node subscribes. Subsequently, the *ckf3D* node saves all the AprilTag positions in two JSON files: one containing the CKF estimates and the other containing the *apriltag_ros* package estimates.

For simulating the TB3, we launch *continuous_detection.launch*, the autonomy stack discussed, and one of the files inside the *TurtleBot3_gazebo* package on the remote PC, such as *TurtleBot3_houseAT.launch*. Each launch file simulates a different gazebo environment. Launch files ending with *AT.launch* are gazebo environments that

contain AprilTags, which allows a complete simulation of our solution.

III. ENVIRONMENT

A. Gazebo World

For the experimental setup, we created two Gazebo [4] environments to simulate a search and rescue mission with AprilTags serving as victims. A total of 12 AprilTags were set up, each measuring around $100\text{mm} \times 100\text{mm}$. Each tag has a unique ID between 0-11 and belongs to the family of *April tag 36h_11* [5]. The reference worlds were *turtlebot3_house* as shown in Fig. 1(a) and *turtlebot3_world* as shown in Fig. 1(b). A third Gazebo environment (Fig. 1(c)) was created to compare *explore_lite* package and our exploration implementation.

B. Real World

For the experimental setup in the real world, we attempted to set up the arena using cardboard and placed 8 AprilTags belonging to the *April tag 36h_11* family, each with a unique ID between 0-7. The maze setup is shown in Fig. 2.



Fig. 2. Example real-world arena setup. Please note that the final setup is larger and includes more AprilTags, as depicted in the demo videos.

IV. METHODOLOGY

A. Proposed Solution

As outlined in Section II, our initial setup utilized existing packages such as *explore_lite* [6] for frontier explo-

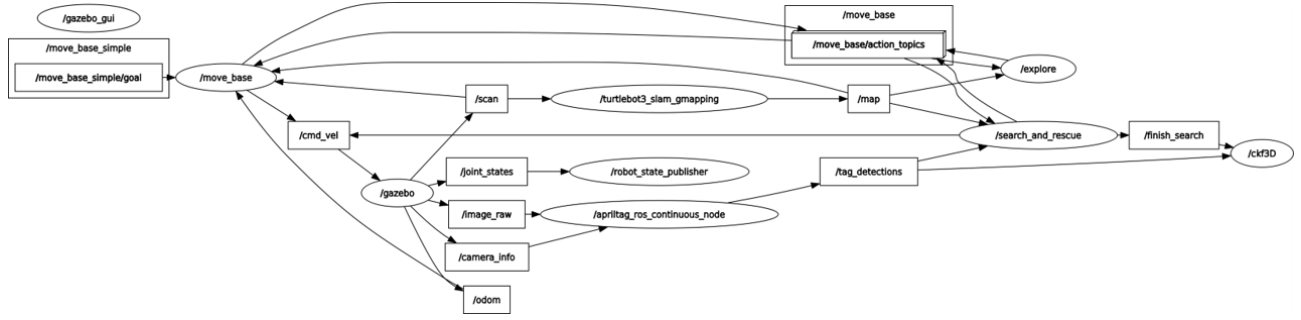


Fig. 3. Nodes and topics in our solution

ration, the `move_base` [7] action server for navigation, `apriltag_ros` [5] for AprilTag detection, and `gmapping` [8] for SLAM. Additionally as described in Section II, custom nodes such as `ckf3D` were employed for recursive Bayesian estimation of AprilTag position, and `search_and_rescue` ensured the entire map was searched again to detect all the AprilTags.

To navigate the TB3 through the environment while effectively mapping and detecting AprilTags, we integrated the ROS package `explore_lite`. Utilizing frontier-based exploration, this package aided in generating a comprehensive map. However, a key challenge arose due to potential oversight of AprilTags during exploration, attributed to the limited field-of-view of the Raspberry Pi Camera. Our attempt to address this involved adjusting parameters within `gmapping`, specifically the minimum and maximum scan distances, under the assumption that these adjustments would enable simultaneous exploration for frontiers and AprilTag detection. However, this approach proved inefficient in achieving the desired outcome.

With `explore_lite` operating as a greedy frontier-based exploration package, prioritization of large frontiers occurs. The influence of frontier size on the frontier cost is determined by parameters such as gain scale and potential scale. By elevating the potential scale to a value exceeding its default of 3, we aimed to prevent the TurtleBot from diverting from its path to pursue distant, larger frontiers, thereby promoting completion of current paths before exploration of new ones.

Furthermore, setting a minimum frontier size, particularly when utilizing `explore_lite` alongside a SLAM-generated map, accounts for potential SLAM mapping oversights, especially in obstacle detection. This parameter specification prevents pursuit of frontiers that may be inaccessible or impractical to explore.

In our solution, we modified only the minimum frontier size while keeping other parameters default. These baseline solutions were integrated with our custom ROS nodes `ckf3D` and `search_and_rescue`. The ROS communication network, illustrated in Fig. 3, depicts interactions among all the nodes and topics pertinent to the autonomous search and rescue stack. This network facilitates seamless communication among software components within the ROS framework [9].

B. Exploration

The `explore_lite` package is used for real world demo. It's a simple frontier-based exploration algorithm that picks the lowest score frontier centroid as the exploration goal. A frontier's cost is computed as a weighted difference between the robot's distance to its closest frontier cell and the frontier's size (number of frontier cells).

It is able to achieve exploration, but has some limitations:

- The chosen exploration goal's orientation is left as zero, which is not optimal for discovering unknown areas, especially if the sensor has limited FOV sensor, such as camera.
- The chosen exploration goal might be unsafe or untraversable: A frontier's centroid may be located in unknown or occupied space. This can cause `explore_lite` to be stuck until a lack of progress timeout is reached, then the node blacklists that goal.
- The frontier detection is not efficient: Every iteration starts anew and performs an entire map search for frontier cells, starting from the closest free cell. This isn't noticeable in small to medium sized maps as the planner frequency is 0.33 Hz, but it's expected to be slow for larger maps.

We implemented a more efficient and effective exploration algorithm that combines ideas from both frontiers and next best view exploration methods.

The exploration algorithm consist of two parts: a) frontier detection and b) computation of exploration goal. The frontier detection is responsible for detecting new frontier cells and forming them into frontiers. This runs continuously in the background. The computation of exploration goal is responsible to return an exploration goal based on the detected frontiers. This runs whenever the associated service is requested.

Frontier Detection: The frontier detection implements the algorithm Expanding Wavefront Frontier Detection (pseudo-code in [10]). It is a BFS algorithm that traverses free cells on the map to detect new frontier cells, as shown in Fig. 4. Each free cell inserted into the queue is marked as visited and evaluated to ascertain its status as a frontier cell. If it's a

free cell and it has at least one unknown neighbor, then it's considered to be a frontier cell. Every detected frontier cell is then stored inside a data structure that only contains frontier cells.

The efficiency of this algorithm comes from these two ideas:

- The visited state is incorporated as an additional layer to the occupancy grid map, so it persists across iterations. Hence, at each iteration, the algorithm only need to search the new, unexplored free cells that are uncovered due to new observations.
- The BFS doesn't naively start searching from robot's current position for unexplored free cells. It starts from frontier cells in the previous timestep F_{t-1} within the observation area A_t . It exploits the fact that new frontier cells can only appear within the region the robot observes and beyond the past frontier cells.




Fig. 4. Expanding Wavefront Frontier Detection. Blue cells are frontier cells from previous timestep. Red cells are new, unexplored free cells that are traversed and evaluated as potential frontier cells. Figure from [11].

The algorithm requires constant insertion and removal of frontier cells, as well as a quick range query of frontiers cells within the observation area. For this reason, we used a RTree data structure to store frontier cells, which is a spatial indexing structure that supports quick range query, similar to KD-tree, but it additionally supports dynamic insertion and removal with self-balancing. We use a bounding box of width and height of the `OccupancyGridUpdate` message from the global costmap as the observation area.

Once the new set of frontier cells F_t has been determined, we perform another search to form frontiers from connected frontier cells. Our implementation of frontiers stores centroid's position and size (number of frontier cells).

Computation of Exploration Goal: As discussed, choosing a position among frontier centroids can have several limitations. For this reason, we used ideas from next best view planners to sample around free space for an optimal exploration goal.

The algorithm for choosing an exploration goal is as follows:

- 1) Sample M points around each frontier's centroid in the free space (Fig. 5).
- 2) Choose an optimal orientation for each sampled point that maximizes the information gain.

- 3) Compute the potential gain for each pose.
- 4) Choose the pose with highest potential gain as the exploration goal (Fig. 5).

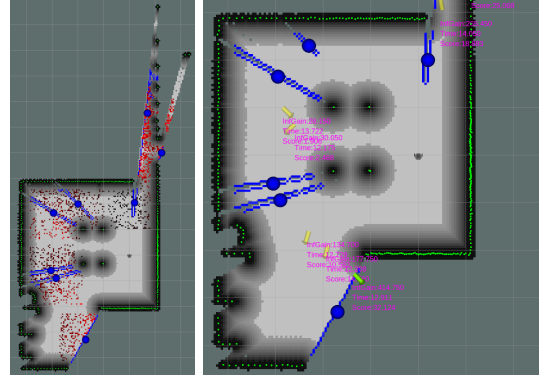


Fig. 5. Computation of exploration goal. a) Sampling process b) Potential gain. The yellow arrows indicate best pose near each frontier, and the green arrow indicate the best pose overall.

We compute the total unknown area that the robot would discover as the information gain, similar to [12]. The information gain and the orientation optimization are implemented according to the next algorithm [13]:

- 1) Discretize 360° into rays.
- 2) For each ray i , count the number of unknowns cells N_i . The information gain for this ray is computed as the length of unknown cells: $I_i = N_i \cdot \text{map resolution}$
- 3) For each orientation θ , the information gain is computed by summing up each ray's information gain in the sensor's field of view:

$$I_\theta = \sum_{i \in \text{sensor fov}} I_i \quad (1)$$

- 4) Choose orientation θ^* that maximizes information gain.

The potential gain $G(X)$ for each pose $X = \{\mathbf{p}, \theta\}$ is computed using information gain and time cost. Let the robot's current pose be denoted as $X_r = \{\mathbf{p}_r, \theta_r\}$ and its maximum velocities be $\{\omega_{max}, v_{max}\}$, the potential cost is given by:

$$G(X) = \frac{I(X)}{T(X)} \quad (2)$$

$$T(X) = \max \left(\frac{\theta - \theta_r}{\omega_{max}}, \frac{\|\mathbf{p} - \mathbf{p}_r\|}{v_{max}} \right) \quad (3)$$

The time cost $T(X)$ is a simple estimate that assumes there is a straight, collision-free path toward the given pose, and the robot is moving with maximum velocities. The information gain $I(X)$ is computed as mentioned. When scoring sampled points, we use their optimal orientation.

C. AprilTag Pose Estimation

Qualitatively, we observed that the `apriltag_ros` package displayed estimator bias in estimating the AprilTag position.

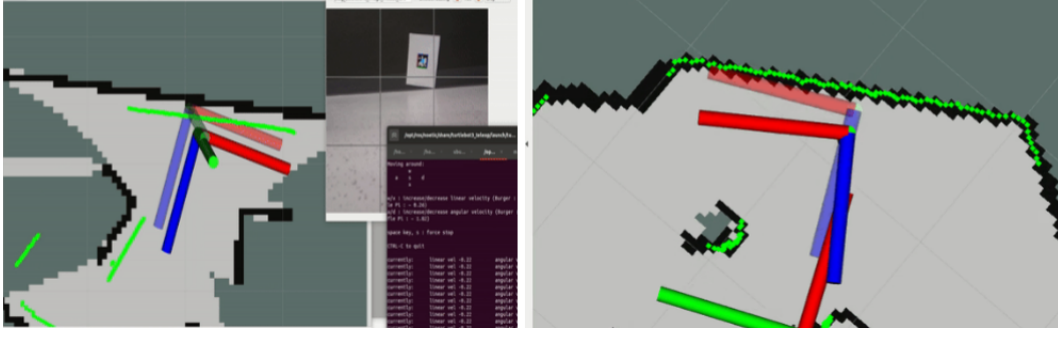


Fig. 6. The lightly shaded axes represent the correct CKF estimator for the AprilTag position, while the darker shaded axes indicate the incorrect apriltag_ros estimator, which should align flush against the wall in RVIZ. The left figure is the real world arena, whereas the right figure is the Gazebo TB3 World arena (using the virtual camera software plugin)




Fig. 7. The geometry of the measurement model when using the World Coordinate Frame and the TB3 Coordinate Frame

As the distance between the TB3 and an AprilTag increased, the estimated AprilTag position drifted farther from the true position in the real world, as observed in RVIZ. To address this issue and other concerns outlined in Section V, we implemented a recursive Bayesian estimator, specifically the CKF [3], for the position estimate of each AprilTag in the World frame.

We utilize the `apriltag_ros` package's relative transformation between the Raspberry Pi camera (`raspicam`) and the AprilTag to construct the measurement model for the 3D AprilTag position, considering range (ρ), bearing (θ_B), and elevation (θ_E) (see Fig. 7 for more details). Leveraging the ROS `tf` package, we can transform the AprilTag and TB3 positions into the World coordinate frame. Since the AprilTags remain stationary, we assume the transition function for the AprilTag is an identity with zero noise. Thus, we may describe the state-space conditional probabilities required for the CKF AprilTag3 position estimator as:

$$p(X_{AT}^{(k)} | X_{AT}^{(k-1)}) = X_{AT}^{(k-1)} + \mathcal{N}(0, 0) \quad (4)$$

$$p(z^{(k)} | X_{AT}^{(k)}; X_{TB3}^{(k)}) = h(X_{AT}^{(k)}; X_{TB3}^{(k)}) + \mathcal{N}(0, R^{(k)}) \quad (5)$$

where we define measurement mean h and measurement noise

covariance $R^{(k)}$ as

$$h(\chi_{AT}; \chi_{TB3}) = \begin{bmatrix} \theta_B \\ \theta_E \\ \rho \end{bmatrix} = \begin{bmatrix} \tan^{-1} \left(\frac{y_{AT} - y_{TB3}}{x_{AT} - x_{TB3}} \right) - \theta_{TB3} \\ \cos^{-1} \frac{\Delta Z}{\sqrt{\Delta X^2 + \Delta Y^2 + \Delta Z^2}} \\ \rho \end{bmatrix} \quad (6)$$

$$R^{(k)} = \begin{bmatrix} 0.05 & 0 & 0 \\ 0 & \frac{\pi}{20} & 0 \\ 0 & 0 & \frac{\pi}{20} \end{bmatrix} \left(\rho^{(k)} \right)^4 \quad (7)$$

In the TB3 coordinate frame, θ_{TB3} represents the positive angle from the x-axis to the AprilTag, while ΔX , ΔY , and ΔZ denote the position of the AprilTag. The geometry of the TB3 and the AprilTag to describe the measurement model is shown in Fig. 7

As the distance between the TB3 and the AprilTag increases, the measurement model covariance matrix also increases, reducing the impact of erroneous measurements on the final position estimate. Additionally, employing repeated measurements helps resolve occlusion issues, as earlier (but not necessarily the last) measurements may have zero occlusion.

D. Search and Rescue

To ensure all AprilTags are detected in the environment, our custom node `search_and_rescue` waits for the `explore_lite` package to finish exploring the environment. Once the `exploration_flag` is received, indicating the completion of the exploration phase and the start of the search and rescue operation, the approach utilizes a grid-based decomposition of the environment represented by the occupancy grid map.

The path planning process includes the following key steps:

- 1) **Inflated Occupancy Grid:** The occupancy grid map undergoes inflation to accommodate the robot's physical dimensions, ensuring obstacles are buffered by a safety margin during path planning. **Free Space Identification:** The inflated map is analyzed to identify unblocked areas, which constitute the free space accessible to the robot.

- 2) Grid Division and Sampling: The free space is divided into a grid structure, with the number of grid cells determined based on the estimated free space area $|grid\ cells| = \sqrt[3]{A_{free}}$. Each cell is iterated upon, and sampling points are generated within its boundaries.
- 3) Valid Point Selection: Sampled points are evaluated against the original occupancy grid map to confirm they reside in free space. If a sampled point falls within an obstacle, the algorithm attempts to find a valid point within the same cell through additional random sampling.
- 4) World Coordinates and Visualization: Valid points, confirmed to be free of obstacles, are transformed from their grid coordinates into world coordinates. These coordinates are stored for navigation purposes and may be visualized using markers for debugging or monitoring.

This grid-based path planning approach efficiently explores the free space in the map, ensuring comprehensive coverage while prioritizing robot safety through obstacle avoidance. The generated set of goal points serves as the foundation for the robot's search path, guiding it towards potential locations of AprilTags during the search and rescue operation. The final result of the sampled points in the occupancy is shown in Fig. 8




Fig. 8. Sampled Points in the occupancy grid

Once all the points are sampled, the node subscribes to the `/tag_detections` topic to continuously receive information about detected AprilTags. Upon receiving data, it parses the information to extract the IDs of the detected AprilTags, storing them in a dictionary. The system continually monitors this dictionary to determine whether all pre-configured AprilTags have been successfully detected.

The sampled points in the queue are sent to the `move_base` action server, responsible for planning a path to the goal and controlling the robot's movements to reach the desired position. Upon reaching each goal point, the robot performs a 360° rotation for 4 seconds at the sampled goal point. Once the rotation is complete, the next sampled point is taken from the queue and sent to the `move_base` action server. This process ensures all AprilTags are detected in the environment. Once all AprilTags are detected in the dictionary, the TB3 ceases exploration.

V. RESULTS

A. Survivor Localization

Due to the challenge of acquiring actual ground truth positions of AprilTags in real-life arenas, we utilize Gazebo to simulate arenas, deliberately placing 11 AprilTags across the map. The ground truth positions of AprilTags can be extracted from the launch files of the Gazebo simulated arenas. We specifically examine two simulated arenas: the TB3 World and the House World.

1) *TurtleBot3 World*: The tb3 World features numerous pillars located in the center of the arena, resulting in occlusion of the AprilTags within the camera's field of view. We hypothesize that this occlusion adversely affects the perspective-n-point transformations used to estimate the relative positions of AprilTags from the raspicam. Consequently, the final measurements from the `april_tag` ROS node may result in inaccurate position estimates of the AprilTags. This issue is exacerbated as the distance between the AprilTag and the tb3 increases. A visual of the `apriltag_ros` estimator, CKF estimator, and ground truth AprilTag position is shown in Fig. 9




Fig. 9. Visualize the ground truth positions, based on the TurtleBot3 World SDF launch file, of 11 AprilTags in the Gazebo TurtleBot3 World. Compare these GT positions with the estimated positions from the CKF and `apriltag_ros`.

2) *House*: The House contains open, spacious rooms with minimal obstacles, resulting in little occlusion of the AprilTags within the camera's field of view. However, the absence of occlusion allows the `apriltag_ros` package to detect AprilTags from greater distances, leading to more erroneous measurements compared to the tb3 World arena. We have observed a direct bias estimation error that is proportional to the distance. A visual of the `apriltag_ros` estimator, CKF estimator, and ground truth AprilTag position is shown in Fig. 10

We observed that the CKF mitigated most of the issues described above. The CKF estimate, which uses many measurements of the same AprilTag, resulted in an improved Mean Squared Error (MSE) for the AprilTag position estimates (Table I).

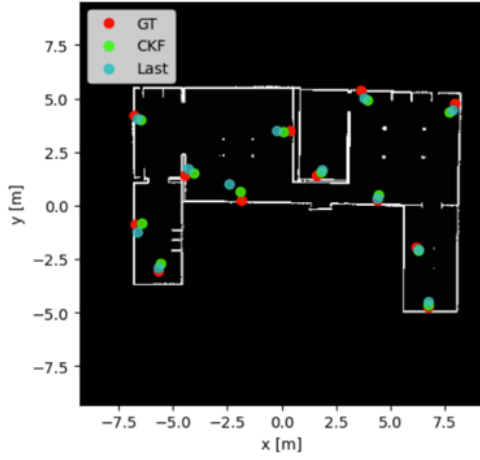


Fig. 10. Visualize the ground truth positions, based on the TurtleBot3 House SDF launch file, of 11 AprilTags in the Gazebo TurtleBot3 House. Compare these GT positions with the estimated positions from the CKF and apriltag_ros.

Tag id	World		House	
	CKF	Last	CKF	Last
Tag0	0.21	0.47	0.32	0.31
Tag1	0.17	0.34	0.34	0.36
Tag2	0.10	0.15	0.20	0.34
Tag3	0.32	0.46	0.49	0.26
Tag4	0.34	0.33	0.31	0.11
Tag5	0.14	0.18	0.11	0.32
Tag6	0.10	0.03	0.22	0.27
Tag7	0.10	0.24	0.34	0.37
Tag8	0.05	0.24	0.28	0.42
Tag9	0.07	0.26	0.28	1.2
Tag10	0.19	0.46	0.37	0.18
Tag11	0.06	0.04	0.33	0.23
Avg	0.15	0.27	0.30	0.36
P-Val	0.02		0.22	

TABLE I

COMPARING THE MSE (IN METERS) OF THE APRILTAG_ROS PACKAGE AGAINST THE CKF FOR THE APRILTAG POSITION ESTIMATE. THE LAST TWO ROWS ARE THE AVERAGE MSE ACROSS ALL APRILTAGS, AND THE UNEQUAL VARIANCE TWO-SAMPLE WELCH T-TEST.


We utilize a two-sample Welch's t-test to determine whether the CKF reduction in MSE is statistically significantly, relative to the apriltag_ros package estimator. The Null Hypothesis is the CKF MSE is larger than the last measurement MSE. For the tb3 World arena, we observe that the CKF indeed significantly reduces the MSE with a P-value of 0.02 (Table I). However, this reduction is not statistically significant for the House area with a P-value of 0.22 (Table I). We hypothesize that with more Monte Carlo trials (and CKF hyperparameter optimization), the CKF would demonstrate statistical significance in reducing the MSE for both arenas.

B. Survivors Detected


We demonstrated that our search and rescue algorithm almost always detects and localizes the AprilTags. However, there were cases where the tb3 failed to detect a single AprilTag (out of 8 AprilTags).

1) *Gazebo Arenas*: In the simulated Gazebo arenas, the tb3 consistently detected and localized all 11 placed AprilTags across three Monte Carlo trials. Examples of successful search and rescue missions in the simulated arena are depicted in Figure Fig. 9 and Figure Fig. 10.

2) *Real World Arena*: In our real-world arena (Fig. 2), we conducted six search and rescue missions to locate eight scattered AprilTags. Out of the six Monte Carlo trials, two successfully found all eight AprilTags, while four found all but one AprilTag. Upon post-analysis of the unsuccessful Monte Carlo trials, we found that the same AprilTag was consistently not detected. This was attributed to the AprilTag being positioned outside the field of view of the tb3 camera, making it physically impossible for the camera to capture it. An example of a successful search and rescue mission, where all eight AprilTags are detected and localized, is illustrated in Figure Fig. 11(a). Conversely, an unsuccessful search and rescue mission, where only seven out of eight AprilTags were detected and localized, is depicted in Figure Fig. 11(b).



(a) success



(b) semi-success

Fig. 11. (Figure (a) depicts a successful search and rescue mission, where all eight AprilTags are detected and localized. However, Tag1 exhibits poor localization performance, possibly due to the selection of incorrect RVIZ axes. In Figure (b), we observe a semi-successful search and rescue mission, where seven out of eight AprilTags are detected and localized. Upon post-analysis, we discovered that the AprilTag was positioned too high for the camera field of view on the tb3 camera, leading to its failure to detect the AprilTag. The lighter shaded axes are the CKF position estimate, whereas the darker axes are the apriltag_ros last measurement estimate

C. Time to Explore Map

The *explore_lite* package and custom exploration package are compared in two gazebo maps: House(video 1) and Maze(video 2). The simulated LIDAR has a range 10m and a 360° FOV.

The configuration of `explore_lite` for both maps is displayed in Table II. The potential scale affects distance to frontier cost, and the gain scale affects frontier size cost. For House gazebo map, it uses the non-inflated map due to exploration failure with costmap.

Parameter	House	Maze
map	gmapping map	global costmap
potential scale	4.0	4.0
gain scale	1.0	1.0
planner frequency	0.33	0.33

TABLE II
EXPLORE LITE CONFIGURATION IN BOTH MAPS

The configuration (only the most important parameters) of custom exploration (package `exploration_alg`) for both maps is displayed in Table III. The sampling radius represents the square region around each frontier centroid where the algorithm samples points from, and it samples the specified number of samples. The sensor range and field of view affect the computation of information gain and the optimized orientation of goal. Although the turtlebot LIDAR sensor can scan 360° , we have limited it so the goal pose always face forward the unknown region. The move base client requests goal every 3 seconds.

Parameter	House	Maze
map	global costmap	global costmap
sampling radius [m]	1.0	1.0
number of samples	50	50
sensor max range [m]	7.0	7.0
sensor field of view [$^\circ$]	[-70, 70]	[-70, 70]

TABLE III
CUSTOM EXPLORATION CONFIGURATION IN BOTH MAPS

The total exploration time for each map is displayed in Table IV. The custom exploration achieves slightly faster exploration time in both maps.

	explore_lite	custom exploration
House	04:44	03:05
Maze	15:30	14:40

TABLE IV
EXPLORATION TIME [MINUTES] COMPARISON

The `explore_lite` package visualizes frontiers as blue lines and displays its score at the initial cell of each frontier as green spheres. When frontiers are blacklisted due to timeout after a lack of progress, they become red. While it does not directly visualize the chosen goal, its implementation suggests that it selects the frontier centroid.

We observed that `explore_lite` tends to prioritize large frontiers, even if the TB3 is far away. Consequently, exploration of the environment often wastes time by traveling to distant regions before completing exploration of local areas. Sometimes, `explore_lite` selects far frontiers, but due to progress timeout, it abandons these goals and pursues closer frontiers, leading to somewhat erratic behavior that paradoxically improves

overall exploration time. Additionally, there are cases where `explore_lite` fails for unknown reasons.

In contrast, our exploration implementation displays frontiers as blue lines, frontier centroids as blue spheres, best sampled poses near each frontier as yellow arrows, and the chosen exploration goal as a green arrow. Our approach appears to make more reasonable choices in selecting exploration goals, favoring close frontiers with sufficient information gain. However, in certain scenarios, such as in the maze map, our implementation could have explored faster, as evidenced by instances where it chose inefficient farther frontiers for their information gain.

To conclude this section, we highlight the need for considering additional criteria for proper comparison, such as total trajectory length and average information gain. These aspects remain avenues for future work.

VI. LESSONS LEARNED

Through our experiences, we have gained insights from the challenges encountered. Let us delve into the four notable issues we encountered along the way.

A. Drift Issue

Drift in mobile robots refers to the unintentional deviation from a planned path or orientation. This can stem from sensor inaccuracies, uneven floor surfaces, or imbalances in the drive mechanism. It's crucial for precision tasks and affects the reliability of autonomous operations. When the OpenCR board, housing the IMU, was not securely fastened to the TB3, it resulted in odometry readings drifting to infinity, causing the TB3 to believe it strayed beyond the map boundaries.

B. Carpet Issue


Operating on carpeted surfaces presents unique challenges for robots, the thickness and texture of carpets can impede movement and operational efficiency. We discovered that the wheel encoder inaccurately recorded movement for the TB3 when the wheels encountered resistance in the carpet, resulting in suboptimal motion control and incorrect odometry.

C. Incorrect Size of April Tag


We initially set the length of the April Tag to 100 mm, but upon printing, we discovered it was actually 70 mm (refer to Section VI-C(a), where the estimated position is inaccurate). Therefore, we needed to re-measure the actual length to avoid any distortion of the grey box of the QR code during scanning. Figure Section VI-C((b) shows accurate estimation after measuring length of April Tag and camera calibration.

D. Frame Rate Drop

This issue affects the visual processing capabilities of robots, where the frame rate of camera feeds unexpectedly drops. It can result in delayed image processing and reduced image quality, crucially impacting tasks that require real-time data interpretation. In real scenario, for the default resolution we desire 15 frames per second. Sometimes it suddenly drops. To



(a) Incorrect AprilTag Edge Size



(b) Correct AprilTag Edge Size

maintain the frame rate, we had to reduce resolution and do camera calibration.

E. Hardware Problems

Robotics systems are susceptible to hardware failures, ranging from motor failures to sensor malfunctions, which can disrupt entire operations. For instance, our OpenCR board experienced a short circuit, necessitating its replacement. Similarly, our LIDAR encountered issues with room scanning, sometimes failing to receive any scans, prompting us to replace it. These steps involved disassembling the TB3 and reassembling it (and rewiring it) with the new installations.

VII. FUTURE WORK

Our primary focus is on integrating and optimizing our exploration and search and rescue algorithms. Presently, our custom exploration algorithm has only been tested in simulation (Gazebo) and remains disconnected with the search and rescue pipeline. However, we acknowledge the imperative need for these algorithms to collaborate during missions, enabling immediate responses upon detecting victims during exploration. This integration ensures a more efficient and cohesive operation of our robotic system.

Furthermore, we emphasize the importance of extensive testing in both simulated and real-world environments. While our exploration methods have undergone simulation testing with various Gazebo world configurations, real-world scenarios demand further evaluation. By experimenting with different layouts in physical arenas, we can validate the efficacy of our algorithms in practical settings. This holistic approach ensures the robustness and reliability of our system across diverse environments. Additionally, addressing the AprilTag position estimation, the measurement model should either leverage the perspective-n-point model or incorporate the measurement bias error as a corrective function.

ACKNOWLEDGMENT

Thanks to Professor Michael Everett for his undying devotion and passion to robotics. Thanks to Donatello for his hardware failing every day.

REFERENCES

- [1] Robotis. Turtlebot3. [Online]. Available: <https://emanual.robotis.com/docs/en/platform/turtlebot3/overview/>
- [2] ROS. How to calibrate a monocular camera. [Online]. Available: http://wiki.ros.org/camera_calibration/Tutorials/MonocularCalibration
- [3] S. Särkkä and L. Svensson, *Bayesian filtering and smoothing*. Cambridge university press, 2023, vol. 17.
- [4] N. Koenig and A. Howard, "Design and implementation of the player/stage/gazebo project," *Proceedings of the 2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2004)*, pp. 2149–2154, 2004.
- [5] J. Wang and E. Olson, "Apriltag 2: Efficient and robust fiducial detection," in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2016, pp. 4193–4198.
- [6] J. Hörner, "Map-merging for multi-robot system," 2016.
- [7] "ROS Wiki: move_base," http://wiki.ros.org/move_base, accessed: April 2024.
- [8] G. Grisetti, C. Stachniss, and W. Burgard, "Improved Techniques for Grid Mapping With Rao-Blackwellized Particle Filters," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2007, pp. 3928–3933.
- [9] Stanford Artificial Intelligence Laboratory et al. Robotic operating system. [Online]. Available: <https://www.ros.org>
- [10] P. Quin, A. Alempijevic, G. Paul, and D. Liu, "Expanding wavefront frontier detection: An approach for efficiently detecting frontier cells," in *Australasian Conference on Robotics and Automation, ACRA*, 2014.
- [11] P. Quin, D. D. K. Nguyen, T. L. Vu, A. Alempijevic, and G. Paul, "Approaches for efficiently detecting frontier cells in robotics exploration," *Frontiers in Robotics and AI*, vol. 8, p. 616470, 2021.
- [12] H. González-Baños and J.-C. Latombe, "Navigation strategies for exploring indoor environments," *I. J. Robotic Res.*, vol. 21, pp. 829–848, 10 2002.
- [13] V. M. Respal, D. Devitt, R. Fedorenko, and A. Klimchik, "Fast sampling-based next-best-view exploration algorithm for a mav," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 89–95.