

Robotic Exploration through Semantic Topometric Mapping

Scott Fredriksson, Akshit Saradagi and George Nikolakopoulos

Abstract—In this article, we introduce a novel strategy for robotic exploration in unknown environments using a semantic topometric map. As it will be presented, the semantic topometric map is generated by segmenting the grid map of the currently explored parts of the environment into regions, such as intersections, pathways, dead-ends, and unexplored frontiers, which constitute the structural semantics of an environment. The proposed exploration strategy leverages metric information of the frontier, such as distance and angle to the frontier, similar to existing frameworks, with the key difference being the additional utilization of structural semantic information, such as properties of the intersections leading to frontiers. The algorithm for generating semantic topometric mapping utilized by the proposed method is lightweight, resulting in the method's online execution being both rapid and computationally efficient. Moreover, the proposed framework can be applied to both structured and unstructured indoor and outdoor environments, which enhances the versatility of the proposed exploration algorithm. We validate our exploration strategy and demonstrate the utility of structural semantics in exploration in two complex indoor environments by utilizing a Turtlebot3 as the robotic agent. Compared to traditional frontier-based methods, our findings indicate that the proposed approach leads to faster exploration and requires less computation time.

I. INTRODUCTION

Robot exploration, often referred to as Active Simultaneous Localization and Mapping (ASLAM), is an automated version of the Simultaneous Localization and Mapping (SLAM) process [1]. ASLAM holds significant importance in mobile robotics due to several reasons. Firstly, it enables a robot to autonomously construct its own map of the operating area. Secondly, it becomes particularly essential in situations where manual control of the robot is impractical, such as in subterranean environments or during search and rescue missions.

The autonomous exploration problem is, in general, divided into three distinct components: pose identification, optimal goal selection, and navigation [2]. Pose detection involves identifying potential locations on the map that the robot can visit to acquire new information about its environment. Optimal goal selection, as the name suggests, is the process of choosing the pose that maximizes both current and future information gain. Lastly, the navigation component safely directs the robot to the chosen goal position.

The majority of the methods in ASLAM focus on the concept of frontiers, which are collections of unoccupied cells on a map adjacent to unexplored regions. Conventional techniques involve identifying these frontiers by locating

their boundaries on the map [3]–[5]. However, this often necessitates scanning large parts of the map, which can be computationally intensive, especially as the map expands in size. An alternative common approach, which is less computationally demanding, is to utilize Rapidly-exploring Random Trees (RRTs) to detect frontiers [6], [7]. A closely related approach to RRTs involves the utilization of a Voronoi graph [8] to direct the exploration, as demonstrated in [9].

Frontier-based exploration has limitations due to its simplistic view of the world as comprising just the regions in between explored and unexplored areas of the map. Therefore, in this article, instead of attempting to address the exploration problem from a frontier-centric perspective, we look to explore the problem from a more semantically meaningful perspective. By representing the world as a collection of intersections, dead ends, and pathways, some of which will lead to more unexplored regions on the map, we leverage this structural semantic information for exploration.

A. Related work

Using semantic information in exploration is a relatively new idea that is likely going to become more popular in the future due to advancements in the fields of semantic mapping [10] and scene graphs [11]. A notable instance of this approach is highlighted in the research of [12], where the explored map was converted to a topometric map divided into semantic "places" that subsequently aid in determining optimal goal positions. However, this approach poses limitations since it presumes that all walls in the map are oriented at 90° angles relative to each other, which limits the types of environments in which the framework can operate. Another example of the use of semantic information utilized in exploration can be found in [13], where semantic labels are integrated into an OcTree map to enhance the exploration process. Their approach uses detailed semantic labels, such as "tree," "building," "grass," and so forth. Similarly, the work in [14] integrates semantic objects into frontiers to locate objects in home environments quickly. In contrast, our proposed framework utilizes fundamental semantic information, emphasizing elements like pathways and intersections.

Another approach, similar to utilizing semantic information for exploration, is the integration of machine learning to enhance robotic exploration. The study presented in [15] employs a generative neural network to predict unmapped areas, thereby optimizing goal selection during the exploration phase. The work in [16] used self-supervised machine learning to do information gain prediction on a 3D

Scott Fredriksson is the corresponding author (email : scofre@ltu.se). The authors are with the Robotics and AI group, in the Department of Computer Science, Electrical and Space Engineering at Luleå University of Technology, Sweden.




Fig. 1: An illustration of the steps taken by the proposed method to explore an unknown environment. The proposed method processes the occupancy map of the currently explored environment (Fig. 1a)), to find the optimal goal location for further exploration (Fig. 1d)), the path to which is smoothly tracked by the robot using modified potential fields (Fig. 1e)).

occupancy grid and detection of semantic objects to improve the exploration process for micro aerial vehicles.

In contrast to the related literature, in this article, we introduce a novel approach to autonomous robot exploration that utilizes structural semantic information through semantic topometric mapping.

B. Contributions

Firstly, we introduce a fundamentally novel and distinct approach to autonomous robotic exploration, wherein structural semantic information contained in the currently explored map, such as properties of intersections, pathways, dead ends, and pathways leading to frontiers, are incorporated to enhance the exploration process. Such semantic information has not been harnessed so far in the literature relating to robotic exploration to the best of the author’s knowledge. The main advantages of the proposed method, when compared to e.g. [3]–[7], stems from the pre-processing stage of the method, where the already explored map is segmented to derive structural semantics. This stage supplies the necessary information for all three stages of exploration: pose identification (unexplored frontiers), optimal goal selection (incorporation of properties of intersections leading to a frontier), and efficient navigation (paths derived from the global skeleton map). In the presented approach, the necessary structural and topological properties are derived from a map produced by LiDAR measurements, as against visual semantics in [13] that require an additional sensor. Finally, the proposed framework is comprehensively compared to an existing frontier-based exploration framework [4] through real-world experiments, which showed that the proposed method completes exploration faster than the framework in [4] while requiring significantly lower computational times for map generation (Figure 9), map exploration (Figure 11) and global navigation (Figure 10).

II. METHODOLOGY

In this Section, we present details of the autonomous robotic exploration strategy proposed in this article, with the overall concept illustrated in Figure 1.

A. Pose Identification

Similar to other exploration frameworks, the proposed method utilizes frontiers as potential exploration poses. The

key difference that sets our proposed approach apart is that the frontiers are derived as part of a semantic topometric map produced by the framework introduced in [17]. This framework processes the currently explored 2D grid-based map, \mathcal{M} , originating from a SLAM solution to yield a semantic topometric map \mathcal{S} . Figure 2 presents an illustrative example of the semantic segmentation performed by the framework proposed in [17]. A topometric map is a hybrid representation, bridging the gap between a grid-based metric map and a topological map [18]. The semantic map created by the framework comprises a set of semantic areas \mathcal{A} , which correspond to various semantic types such as intersection, pathway, dead-end, and paths that lead to unexplored areas. Additionally, the framework outputs a set of paths represented as a skeleton map of the environment, which can aid a global navigator in swiftly and efficiently navigating the map. Among the semantic areas are N_f pathways leading to frontiers, collected into $\mathcal{A}_F \subset \mathcal{A}$. To each $A_i^F \in \mathcal{A}_F$, $i \in \{1, \dots, N_f\}$, there is a global path P_i that connects the robot’s current pose to the associated unexplored frontier F_i . The set $\mathcal{F} = \{F_i\}$ is then the collection of all target poses for robot exploration.

B. Optimal Goal Selection

A common approach for goal selection, adopted by many exploration frameworks, is the use of the metric information of the frontier, distance, and angles to determine the optimal goal position $F^* \in \mathcal{F}$. This is often done without taking the map into consideration, primarily because it is computationally expensive to calculate paths to every frontier on

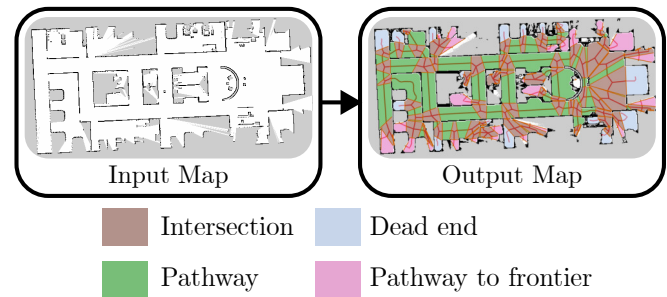


Fig. 2: An example of a semantic topological map produced by the method in [17].

the map. In contrast, when utilizing the semantic topometric map, paths can be found to all possible poses relatively computationally cheaply, as described in subsection II-C. Therefore, the distance cost is determined by the length d of the path from the robot to each pose, and the angle cost v is determined by the difference between the robot's current pose and the pose at the start of the global path to the frontier, illustrated in Figure 3.

In addition to the length d and angle cost v , semantic information from the topometric map is employed in the proposed method to determine the optimal exploration goal F^* . The first consideration is the length $p_l(F_i)$ of the path between a frontier F_i and the nearest intersection denoted as I_i^F . This is motivated by the observation that paths leading to unexplored regions, which don't lead anywhere (such as walls hidden behind other walls at intersections), are typically very short. The proposed method also uses the properties of the intersection connected to the path leading to the unexplored area I_i^F . Specifically: 1) The number of openings, $O(I_i^F)$, in the intersection I_i^F , with the assumption that an intersection with numerous openings is significant in the overall map, and hence, exploring paths connected to it will yield higher information gain than those linked to intersections with fewer openings and 2) The number of pathways $P_u(I_i^F)$, connected to the intersection directly leading to a frontier, as a higher number of connected pathways suggests higher future information gain, resulting in lesser backtracking for a robot.

The equation below consolidates all the above cost and gain considerations and presents a single metric for each frontier F_i :

$$C(F_i) = C_{metric} - G_{semantic} \quad (1)$$

where

$$C_{metric} = \frac{d}{L_{path}} + w_v \cdot v \quad (2)$$

$$G_{semantic} = w_p \cdot p_l(F_i) + w_I \cdot (O(I_i^F) + P_u(I_i^F)).$$

The terms in Equations (1) and (2) are defined as follows: C_{metric} is the cost of the path using the metric properties of the path. $G_{semantic}$ measures the semantic value gain of the path based on the semantic properties along the path leading to the frontier. Finally, w_v , w_p , and w_I are weights set by the user to determine the importance of the various factors. The weighting parameters are tuned based on the following rationale:

1) *Tuning of w_v* : A higher value of w_v decreases the likelihood that the robot will turn around, thereby committing the robot to the direction in which it is currently headed. This is advantageous for avoiding local cost minima, but an excessively high value of w_v may cause the robot to make poor decisions, leading it to take unnecessary detours to avoid turning around.

2) *Tuning of w_p and w_I* : Low values for w_p and w_I are preferred to ensure that the robot does not bypass nearby frontiers in favor of more distant but complex frontiers. However, w_p and w_I should be sufficiently high to allow the




Fig. 3: Illustration of a path leading to a frontier in a semantic topometric map, highlighted in blue, for a specific scenario. The figure also shows the different semantic entities used in the optimal selection process presented in Section II-B. The intersection connected to the target frontier has four openings, and only one of its pathways leads directly to a frontier.

robot to choose a nearby frontier with a higher potential for information gain when faced with multiple nearby options.

The variable L_{path} represents the average length of all paths between intersections and serves as a normalizing factor. In more complex environments where the distances between intersections are small, the weights w_v , w_p , and w_I become less critical because the distance d is generally lower. Conversely, in larger maps with greater distances between intersections, both the distance d and the weights are likely to be higher. This ensures that the robot's behavior remains consistent regardless of the complexity of the environment. With the cost associated with each frontier F_i as defined in (1), the optimal goal selection reduced to the following optimization problem.

$$F^* = \arg \min_{F_i \in \mathcal{F}} (C(F_i)) \quad (3)$$

The algorithm to find the path P that connects the robot's current pose to F^* , from the skeleton map of the environment, is presented in the next subsection.

C. Navigation

The proposed Algorithm 1 outlines the method used to find the optimal path, utilizing the cost function from Eq. (1). The cost is applied at two points in the algorithm; the angle cost $w_v \cdot v$ is calculated at the start of each path (at line 3), and the remaining costs are calculated after finding the complete path (at line 9). To optimize the search process, the algorithm incorporates one additional condition. The condition specifies that if a path reaches an area that another path has already reached with a lower score, the current path is discontinued (The current value of each area is tracked by the A_{cost} array).

Utilizing pre-calculated global paths in the semantic topometric map is computationally efficient. However, as these paths are general global paths generated without any consideration of the robot's current position, they are typically suboptimal. This is true not only in terms of the shortest point-to-point maneuver but also with respect to smoothness and length (see Figure 1b)). To mitigate these shortcomings,

Algorithm 1 Method To Find Optimal Path

Input: $firstNode$: The semantic area robot is in.
Initiate: $searchList$: The paths being searched.
 $candidate$: the current best path.
 A_{cost} : array of size A , set to ∞

```

1: for  $N \in get\_neighbour(firstNode)$  do
2:   Compute  $v$  from robot to path to  $N$ 
3:   Add new  $N$  to  $searchList$  with computed cost  $w_v * v$ 
4: end for
5: while  $searchList \neq \emptyset$  do
6:    $c_{path} \leftarrow$  last element of  $searchList$ 
7:   Remove  $c_{path}$  from  $searchList$ 
8:   if  $c_{path} \in \mathcal{F}$  then
9:     if  $cost(c_{path}) < candidate.score$  then
10:      Update  $candidate$ 
11:    end if
12:  end if
13:  for  $N \in get\_neighbour(c_{path}.last)$  do
14:    Compute  $d$  between  $c_{path}.last$  and  $N$ 
15:    If  $A_{cost}[N] < c_{path}.cost + d$  continue
16:     $c_{path} \leftarrow c_{path} + N$ 
17:     $c_{path}.cost \leftarrow c_{path}.cost + d$ 
18:     $A_{cost}[N] \leftarrow c_{path}.cost$ 
19:    Add  $c_{path}$  to  $searchList$ 
20:  end for
21: end while
Output:  $candidate$  : as a path  $P$ 

```

we propose a novel local navigation scheme using a modified version of a potential field.

We consider a robot equipped with a 2D LiDAR sensor that receives a set of N range measurements, denoted as $\mathcal{L} = \{L_1, L_2, \dots, L_i, \dots, L_N\}$, during each update. Let $\mathcal{L}_v = \{L_{v1}, L_{v2}, \dots, L_{vi}, \dots, L_{vN}\}$ represent the angles corresponding to each LiDAR range measurement. The potential field is described by the following equation:

$$\bar{F} = f_g \begin{bmatrix} \cos(G_v) \\ \sin(G_v) \end{bmatrix} - \sum_{i=1}^N \frac{f_o}{L_i} \begin{bmatrix} \cos(L_{vi}) \\ \sin(L_{vi}) \end{bmatrix} \quad (4)$$

In Eq. (4), \bar{F} is the motion vector relative to the robot. The parameters f_g and f_o are tunable factors affecting the potential field. The angle G_v is the direction towards the goal position G , which is a point on the global path that is not obscured by any obstacles and is within a distance of $f_l * L_{\min}$, where $L_{\min} = \min(\mathcal{L})$ and f_l is a scaling factor, as illustrated in Figure 4 that shows the effect of different choices of f_l . The potential field is tuned using the following method: The value f_o is increased until the robot does not collide with walls. The value f_g is incremented until the robot doesn't get stuck in local force minima. The value f_l determines the strictness of the path; a lower value will result in the robot closely following the path but may lead




Fig. 4: Illustration of the goal selection for point G on the path represented as a blue line. The area detectable by the LiDAR, up to a range of $f_l * L_{\min}$ for different values of f_l (f_{l1} , f_{l2} and f_{l3}), is marked by dotted lines.

to suboptimal routes to reach the goal. This is illustrated by the value f_{l1} in Figure 4. Conversely, a higher value of f_l allows the robot to pursue faster, more optimal paths but might bring it closer to walls, increasing the risk of collisions, as illustrated by the value f_{l3} in Figure 4.

III. VALIDATION AND COMPARISON

A. Experimental Setup

The Proposed Method (PM) was benchmarked against an existing framework, *frontier_exploration* (FE) [4]. The latter is a wavefront frontier detector (WFD) [19] based approach, which employs both frontier size and distance metrics for optimal goal selection. Both frameworks were evaluated in two distinct scenarios, utilizing the TurtleBot3 robot platform. The robot navigated through two maze configurations shown in figure 5, that were custom-built for the validation tests. For each maze configuration, 10 successful exploration missions



(a) Maze Configuration 1



(b) Maze Configuration 2

Fig. 5: The two test environments used in the validation and testing.

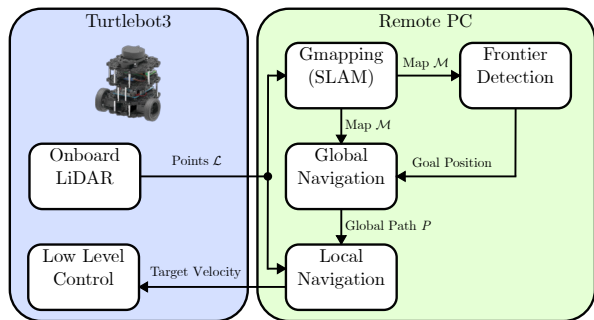


Fig. 6: Robot Software Stack and implementation of the proposed exploration strategy.

were performed with both frameworks. Both frameworks were run until they self-stopped, which is when they deemed that the complete map had been explored (no more frontiers left to explore). The software stack of the robot is detailed in Figure 6. Both frameworks utilize *Gmapping* [20] as their SLAM method. The map resolution was chosen to be 0.05 m. FE uses *Navfn* as its global planner and *DWA* as its local planner, both of which are components of the navigation stack part of the Robot Operating System (ROS) [21]. In contrast, PM uses the semantic topometric map [17] for frontier detection, with global navigation strategies and local navigation presented in subsection II-C. Optimal goal selection, presented in subsection II-B, is performed during the frontier selection process in the FE framework, whereas in PM, it is carried out in the global navigation step. Most of the ROS nodes were executed on a remote PC equipped with an AMD 5850u CPU and running Linux, kernel version 6.4.12.

To evaluate and compare the exploration performance of both methods, we use four metrics: a) the unoccupied area explored over time, b) the time required to complete the exploration process, c) the computational time for frontier detection, and d) the computational time for global navigation. In Tables I, we list the parameters used in the PM. The Turtlebot3 robot platform’s maximum heading speed was limited to 0.15 m/s, and maximum rotation speed was limited to 2 rad/s.

Parameter	w_v	w_p	w_I	f_g	f_o	f_l
Value	4	0.2	0.1	30	0.065	1.5

TABLE I: Values chosen for the tuning parameters of the proposed algorithm for validation. w_v , w_p , and w_I are the parameters used in optimal goal selection. f_g , f_o and f_l are the parameters used in navigation.

B. Experimental Results

The results of the average explored area over time for all ten runs for both PM and FE are depicted in Figure 8 for both test scenarios. The graphs indicate that PM achieves a similar exploration gain to FE at the start of the exploration mission. However, PM explores in a manner




Fig. 7: The maps generated by the run with median exploration time. The starting position of the robot is indicated by a red square, while its final position is denoted by a red star. The path traversed by the robot during its exploration is represented by a purple line. On the map produced by PM the semantic topometric map, generated by the method in [17] is shown as well.

that results in less backtracking towards the end of the mission, as can be seen in Figures 7 and 8. This results in PM continuing to efficiently explore for longer, resulting in a higher average exploration speed and it finishing its exploration mission faster than FE, as is evident in Figures 8 and 11. Another problem that causes the FE method to be slower in exploration time is that it does not check if a frontier is reachable before committing to it, resulting in the robot having to stop its exploration until it can find a frontier it can reach. Regarding computational performance, Figure 9 and 10 show that both the semantic map generation and navigation used by PM are significantly faster than the frontier detection and navigation adopted by FE.

C. Comparison and Discussion

The PM consistently outperformed the FE in terms of exploration speed, time to complete exploration and computational processing time. A significant takeaway from these results is the demonstrated advantage of preprocessing the currently explored map to construct a comprehensive semantic topometric map, rather than merely detecting frontiers, to enable more efficient exploration and navigation. The structural semantics of the explored map hold crucial information for current and future information gain (as demonstrated by the proposed method), which is missed by the state-of-the-art exploration strategies. The noticeable advantages of minimal

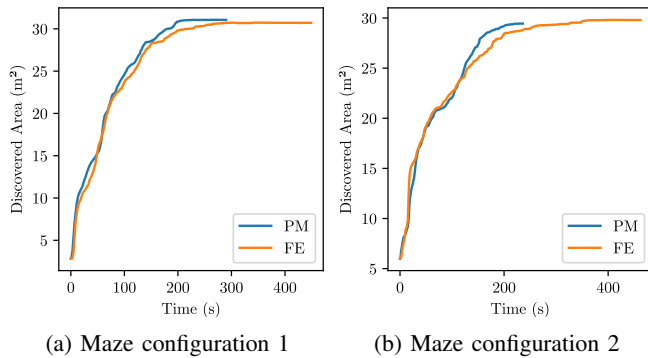


Fig. 8: The explored area of unoccupied space (m^2) over time, averaged over 10 runs.




Fig. 9: Comparison in compute time (ms) between semantic topological map generation using method in [17] for PM and frontier detection in FE

back-tracking and consistent rise in the information gained throughout the exploration mission are due to the utilization of structural semantic information.

As seen in Figure 10, global navigation in the PM is substantially faster than the global planner used by FE. Additionally, the global planner for PM identifies paths to all frontiers, whereas the FE’s planner determines only one path at a time. A critical aspect of global planners is their scalability. The FE’s planner scales with the map’s resolution; in this experiment, the map resolution was 224×224 , which is relatively small for an occupancy map. Thus, on a standard or larger sized map, FE’s navigation performance is expected to deteriorate generally. In contrast, the PM’s global planner scales with the number of intersections, making it more scalable than the FE’s planner, as it scales with the complexity (number of intersections) and not the actual size of the map. An important point of distinction between the semantic-based exploration in PM to existing solutions using semantics is that, while other solutions employ multiple sensors, such as cameras, LiDAR, etc., to detect and classify semantic objects, PM only uses the map generated from SLAM as its input to generate structural semantics. This makes the approach in PM much more fundamental and easier to implement. Another crucial observation is that PM does not utilize any machine learning, which implies that its performance is more reliable in unknown environments

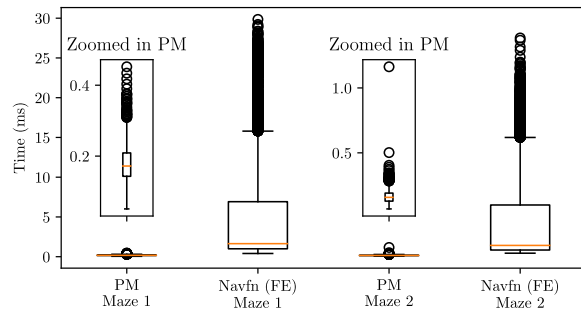


Fig. 10: Comparison in compute time (ms) between global navigation for PM and FE




Fig. 11: Comparison in exploration mission time (s) between PM and FE.

compared to machine learning methods that might face issues in environments that are very dissimilar from their training dataset.

IV. CONCLUSIONS

In this article, we introduced a novel strategy for robotic exploration that utilizes a semantic topometric map. This method departs from traditional frontier-centric approaches by incorporating structural semantic information into the exploration strategy, along with the metric information related to the frontiers. The structural semantics are derived solely from the LiDAR and the occupancy map of the explored parts of the environment by segmenting the map into semantically meaningful areas like intersections, pathways, and dead-ends [17]. The proposed method incorporated into optimal goal selection semantic elements such as the distance of the frontier from the nearest intersection and the number of openings out of the intersection. A modified potential field mechanism was also proposed to dynamically adjust path-following based on environmental contexts. This ensures more efficient navigation to compensate for the suboptimal paths in the global skeleton map generated by [17]. Experimental validation demonstrated the superior performance of the proposed method over the *frontier_exploration* (FE) technique [4] in terms of exploration rate, exploration time, and computational time. In the future, we aim to further investigate the semantic information utilized in the exploration, as well as conduct a more in-depth study on how the different semantic properties influence the exploration process.

REFERENCES

- [1] J. A. Placed, J. Strader, H. Carrillo, N. Atanasov, V. Indelman, L. Carlone, and J. A. Castellanos, "A Survey on Active Simultaneous Localization and Mapping: State of the Art and New Frontiers," *IEEE Transactions on Robotics*, vol. 39, no. 3, pp. 1686–1705, June 2023.
- [2] I. Lluvia, E. Lazkano, and A. Ansuategi, "Active Mapping and Robot Exploration: A Survey," *Sensors*, vol. 21, no. 7, p. 2445, Jan. 2021. [Online]. Available: <https://www.mdpi.com/1424-8220/21/7/2445>
- [3] 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'*, July 1997, pp. 146–151.
- [4] A. Topiwala, P. Inani, and A. Kathpal, "Frontier Based Exploration for Autonomous Robot," Tech. Rep., June 2018, arXiv:1806.03581 [cs] type: article. [Online]. Available: <http://arxiv.org/abs/1806.03581>
- [5] Z. Sun, B. Wu, C.-Z. Xu, S. E. Sarma, J. Yang, and H. Kong, "Frontier Detection and Reachability Analysis for Efficient 2D Graph-SLAM Based Active Exploration," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Oct. 2020, pp. 2051–2058, notissn: 2153-0866.
- [6] H. Umari and S. Mukhopadhyay, "Autonomous robotic exploration based on multiple rapidly-exploring randomized trees," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Sept. 2017, pp. 1396–1402, notissn: 2153-0866.
- [7] D. H. Q. Tran, H.-A. Phan, H. D. Van, T. Van Duong, T. T. Bui, and V. N. T. Thanh, "An Enhanced Sampling-Based Method with Modified Next-Best View Strategy For 2D Autonomous Robot Exploration," in *2023 20th International Joint Conference on Computer Science and Software Engineering (JCSSE)*, June 2023, pp. 225–230, notissn: 2642-6579.
- [8] F. Aurenhammer, "Voronoi Diagrams—a Survey of a Fundamental Geometric Data Structure," vol. 23, no. 3, pp. 345–405.
- [9] L. Li, X. Zuo, H. Peng, F. Yang, H. Zhu, D. Li, J. Liu, F. Su, Y. Liang, and G. Zhou, "Improving Autonomous Exploration Using Reduced Approximated Generalized Voronoi Graphs," *Journal of Intelligent & Robotic Systems*, vol. 99, no. 1, pp. 91–113, July 2020. [Online]. Available: <https://doi.org/10.1007/s10846-019-01119-6>
- [10] I. Kostavelis and A. Gasteratos, "Semantic mapping for mobile robotics tasks: A survey," vol. 66, pp. 86–103. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0921889014003030>
- [11] X. Chang, P. Ren, P. Xu, Z. Li, X. Chen, and A. Hauptmann, "A Comprehensive Survey of Scene Graphs: Generation and Application," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 45, no. 1, pp. 1–26, 2023.
- [12] A. Cowley, C. J. Taylor, and B. Southall, "Rapid multi-robot exploration with topometric maps," in *2011 IEEE International Conference on Robotics and Automation*, May 2011, pp. 1044–1049, notissn: 1050-4729.
- [13] A. Asgharivaskasi and N. Atanasov, "Semantic OcTree Mapping and Shannon Mutual Information Computation for Robot Exploration," *IEEE Transactions on Robotics*, vol. 39, no. 3, pp. 1910–1928, June 2023.
- [14] J. Chen, G. Li, S. Kumar, B. Ghanem, and F. Yu, "How To Not Train Your Dragon: Training-free Embodied Object Goal Navigation with Semantic Frontiers," *Robotics: Science and Systems (RSS) 2023*, May 2023, arXiv:2305.16925 [cs] type: article. [Online]. Available: <http://arxiv.org/abs/2305.16925>
- [15] R. Shrestha, F.-P. Tian, W. Feng, P. Tan, and R. Vaughan, "Learned Map Prediction for Enhanced Mobile Robot Exploration," in *2019 International Conference on Robotics and Automation (ICRA)*, May 2019, pp. 1197–1204, notissn: 2577-087X.
- [16] Y. Tao, Y. Wu, B. Li, F. Cladera, A. Zhou, D. Thakur, and V. Kumar, "SEER: Safe Efficient Exploration for Aerial Robots using Learning to Predict Information Gain," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, May 2023, pp. 1235–1241.
- [17] S. Fredriksson, A. Saradagi, and G. Nikolakopoulos, "Semantic and topological mapping using intersection identification," in *Proc. of the 22nd IFAC World Congress, International Federation of Automatic Control, Yokohama, Japan*, 2023, pp. 9954–9959.
- [18] E. Remolina and B. Kuipers, "Towards a general theory of topological maps," vol. 152, no. 1, pp. 47–104, 2004.
- [19] M. Keidar and G. A. Kaminka, "Robot exploration with fast frontier detection: Theory and experiments," in *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems-Volume 1*, 2012, pp. 113–120.
- [20] G. Grisetti, C. Stachniss, and W. Burgard, "Improved Techniques for Grid Mapping With Rao-Blackwellized Particle Filters," *IEEE Transactions on Robotics*, vol. 23, no. 1, pp. 34–46, Feb. 2007.
- [21] M. Quigley, K. Conley, B. Gerkey, J. Faust, T. Foote, J. Leibs, R. Wheeler, A. Y. Ng, *et al.*, "Ros: an open-source robot operating system," in *ICRA workshop on open source software*, vol. 3, no. 3.2. Kobe, Japan, 2009, p. 5.