

Active Semantic Mapping and Pose Graph Spectral Analysis for Robot Exploration

Rongge Zhang, Haechan Mark Bong, Giovanni Beltrame

Abstract—Exploration in unknown and unstructured environments is a pivotal requirement for robotic applications. A robot’s exploration behavior can be inherently affected by the performance of its Simultaneous Localization and Mapping (SLAM) subsystem, although SLAM and exploration are generally studied separately. In this paper, we formulate exploration as an active mapping problem and extend it with semantic information. We introduce a novel active metric-semantic SLAM approach, leveraging recent research advances in information theory and spectral graph theory: we combine semantic mutual information and the connectivity metrics of the underlying pose graph of the SLAM subsystem. We use the resulting utility function to evaluate different trajectories to select the most favorable strategy during exploration. Exploration and SLAM metrics are analyzed in experiments. Running our algorithm on the Habitat dataset, we show that, while maintaining efficiency close to the state-of-the-art exploration methods, our approach effectively increases the performance of metric-semantic SLAM with a 21% reduction in average map error and a 9% improvement in average semantic classification accuracy.

I. INTRODUCTION

Autonomous exploration and scene understanding constitute pivotal capabilities for advancing robotic systems. Conventionally, exploration relies on geometric maps to facilitate navigation and obstacle avoidance. However, augmenting a robot’s autonomy necessitates a comprehensive understanding of the environment’s semantic information. Metric-semantic Simultaneous Localization and Mapping (SLAM) emerges as a key solution, offering a dual perspective of geometric and semantic environmental perception. This is invaluable for wide applications, from indoor service robots to large-scale outdoor exploration. However, current approaches often employ passive robot motion control during SLAM, limiting the potential for dynamic exploration.

In previous literature on robotic exploration, the information-theoretic approaches are promising solutions, showing the ability to facilitate rapid exploration and minimize map uncertainty. Nevertheless, these approaches often segregate the active mapping problem from localization, i.e., assuming that localization errors have been resolved. This assumption can introduce constraints in practical scenarios, such as compromised path planning due to SLAM errors and incorrect object labeling in semantic maps, causing the robot to pick up wrong objects.

The authors are with the Department of Computer Engineering and Software Engineering, Polytechnique Montréal, Montréal, Canada. {rongge.zhang, haechan.bong, giovanni.beltrame}@polymtl.ca

This work was supported by the National Research Council Canada (NRC).




Fig. 1: Representation of our problem: a robot explores a Habitat environment, building a 3D semantic map using depth sensors and segmentation, and using it for navigation.

Motivated by these gaps, this paper focuses on active metric-semantic SLAM. Our approach extends existing information-theoretic methods for binary [1] or multi-class [2] active mapping by incorporating localization uncertainty into exploration. We maintain a pose graph and utilize a novel quantitative metric of pose graph uncertainty, inspired by recent advances in spectral graph theory. Analyzing the semantic mutual information and underlying pose graph topology allows us to evaluate the maximum gain of different potential trajectories, thus promising higher-quality localization and semantic mapping while maintaining efficient exploration (see Fig. 1 as an illustration). In summary, the main contributions of this paper are:

- 1) A novel approach to active SLAM with online decision-making for autonomous robot exploration;
- 2) A real-time metric-semantic mapping system for active robot perception and scene understanding;
- 3) A hybrid representation integrating semantic mutual information gain and localization uncertainty to guide the robot’s action policy;

We demonstrate the performance of our method through extensive experiments in photorealistic environments, our implementation is publicly available at: https://github.com/BohemianRhapsodyz/semantic_exploration.

II. RELATED WORK

A. Semantic Mapping

Semantic mapping aims to construct semantically annotated 3D maps during robotic operations. Initial approaches

focused on embedding semantic labels into the geometric spatial representation of the map such as voxel, surfel, or mesh, updating them via Bayesian probabilistic fusion upon receiving new observations, typically by SemanticFusion [3] and Voxelbloxx [4]. To achieve higher-level scene understanding, hybrid representations such as scene graphs are introduced, pioneering works include Kimera [5] and Hydra [6]. Recent works like Concept-Graph [7], which implements a multimodal correspondence between language and vision based on Contrastive Language-Image Pre-training (CLIP) [8], can be used for text querying and language interaction navigation. Moreover, neural representations have emerged as a promising frontier in semantic mapping, although they often come with computational overhead and real-time constraints. Examples include Neural Radiance Fields (NeRF) [9] and Gaussian-Splatting-based mapping methods [10].

B. Exploration

Exploration strategies in robotics have significantly evolved since the seminal concept of frontier exploration, different exploration methods have been proposed, ranging from frontier selection to optimization of travel costs. For the application in recent years, GBPlanner [11] was successfully used in the DARPA Subterranean Challenge. TARE planner [12] proposed a hierarchical exploration strategy directly on point cloud maps and has been extended to multi-robot exploration. Among these different methods, information-theoretic exploration aims to achieve the balance between unexplored space and robot state uncertainty. Approaches like Cauchy-Schwarz quadratic mutual information (CSQMI) [13] and fast Shannon mutual information (FSMI) [1] have achieved fast exploration on the binary-class map. Semantic Shannon mutual information (SSMI) [2] followed up by developing a closed-form expression for multi-class mutual information and successfully showing its capabilities for exploration with semantic maps. A common limitation across these methodologies is the assumption of perfect localization, which is impractical in reality. Recent work [14] evaluates the frontier gain and considers localization uncertainty. Through ablation experiments, it elucidates the importance of localization uncertainty in influencing exploration performance. However, it stops short of addressing the navigation challenges this uncertainty introduces.

C. Active SLAM

Active SLAM is usually considered to be the problem of reducing SLAM localization and mapping uncertainty by actively controlling the motion of the robot (or sensors attached to the robot) and is often decomposed into three phases: goal determination, utility computation, and action selection [15]. Carlone et al. [16] propose a utility function for discounting the information gain of an action based on the probability of having good localization and intentionally guiding the robot to revisit previously explored environments to increase potential loop closure constraints. However, their approach is limited to particle-filtered SLAM.

Branching off information theory, the theory of optimal experimental design (TOED) is introduced to evaluate the uncertainty of potential trajectories, with Carrillo et al. [17] showing that D-optimality, which guarantees monotonicity during exploration, is an ideal criterion for TOED and proposing Shannon-Rényi entropy [18] to evaluate varying magnitudes of map and localization uncertainty. Placed et al. [19] present a general relationship between optimality criteria and connectivity indices for active graph-SLAM and use it for fast exploration. However, the map points are marginalized during the construction of the pose graph, the expected gain is calculated as a simple unexplored cell ratio, and the uncertainty of the map is not adequately taken into account. In recent work, Tao et al. [20] propose a 3D active SLAM system for drones by using semantic information for detecting loop closure and estimating the relative pose transformations. Distinguishing from their work, this paper uses semantics for mutual information computation and focuses on the utility function in active SLAM.

D. Spectral Graph Theory

Optimal design criteria, such as D-optimality, offer a scalar representation of the information matrix of the SLAM estimation, effectively quantifying the uncertainty. While these metrics have significantly advanced active SLAM research, evaluating them involves computing a large and dense covariance matrix, which is still expensive, limiting its usage for real-time decision-making. Notably, it has been shown that there is a close relationship between the optimality and topology structure of the underlying pose graph [21]. Specifically, maximizing optimality equates to maximizing the algebraic connectivity of the underlying graph, which can be determined by analyzing the Laplacian spectrum of the topology, thus avoiding the costly computational burden associated with calculating the entire posterior covariance matrix. This has spurred many advancements in active SLAM [22]. Interestingly, beyond active SLAM, the study of spectral graph theory has also been widely applied in other SLAM fields such as rotation averaging and translation estimation [23], pose graph sparsification [24] and collaborative SLAM [25].

III. PRELIMINARIES

A. Occupancy Grid Mapping and Mutual Information

As is standard in information-theoretic exploration, we model the environment in occupancy grid. The map m consists of n grids that can be modeled as $m = \{m_1, m_2, \dots, m_n\}$, where $m_i \in k$ indicates k categories. For the binary mapping, $k \in \{0, 1\}$ for free and occupied cells, respectively. The observation is from a range measurement sensor such as LiDAR or RGB-D camera equipped on the robot. Assume we have an observation Z_t and robot pose X_t at time t , the probability density function (PDF) of map $p_t(m)$ can be updated using Bayesian rule by:

$$p_t(m) \propto p(Z_t | m, X_t) p_{t-1}(m). \quad (1)$$

Mutual information (MI) is introduced and devoted to reducing mapping uncertainty (entropy). Given a new observation, minimizing the map's conditional entropy is equal to maximizing the MI between m and Z_t :

$$\begin{aligned} I(m; Z_t = z_t \mid z_{1:t-1}, x_{1:t-1}) &= \\ H(m \mid z_{1:t-1}, x_{1:t-1}) - H(m \mid Z_t = z_t, z_{1:t-1}, x_{1:t-1}). \end{aligned} \quad (2)$$

For a beam-based sensor with finite range Θ beams observation, the MI between the map and the measurements is the sum of the MI for every beam $\theta \in \Theta$ in a sensor scan at time t :

$$\begin{aligned} I(m; Z_{t,\Theta} \mid Z_{1:t-1,\Theta}) &= \int_{\theta \in \Theta} I(m, Z_{t,\theta}) d\theta \\ &= \int_{\theta \in \Theta} \sum_{m \in k} p(m, Z_{t,\theta} \mid Z_{1:t-1,\theta}) \\ &\quad \times \log \frac{p(m, Z_{t,\theta} \mid Z_{1:t-1,\theta})}{p(m \mid Z_{1:t-1,\theta}) p(Z_{t,\theta} \mid Z_{1:t-1,\theta})} d\theta. \end{aligned} \quad (3)$$

Thus, by accumulating the MI over the independent grids and counting the overall MI, we can evaluate every potential trajectory and select the one that maximizes the information gain for robotic exploration and active mapping.

B. Graph-based SLAM and Optimality Criteria

The SLAM estimation problem can be solved with graph methods, where the node and edge represent the pose and relative motion, respectively. The relative estimation is perturbed with noise under the Gaussian distribution assumption. In this paper, we focus on pose graph optimization, i.e. the landmark or map points are marginalized (the mapping uncertainty will still be considered in the semantic mapping part of this work). In short, the pose graph optimization problem can be treated as a least-square solving problem:

$$\min_x f(x) = \min_x \frac{1}{2} \sum_{i,j \in \varepsilon} e_{ij}^T(x) \Sigma_{ij}^{-1} e_{ij}(x). \quad (4)$$

where x is the state (robot pose), ε is the set of all edges, $e_{ij}(x)$ is the error to be minimized and Σ_{ij} is the covariance matrix. Several optimality criteria can be applied to SLAM according to the theory of optimal experimental design. A-optimality, T-optimality, E-optimality, and D-optimality [17] are mostly used criteria that optimize different properties of the covariance matrix in pose graph optimization. Among them, only D-optimality captures the global variance uncertainty and keeps the monotonicity [17], the representation of D-optimality is:

$$D-opt = \exp\left(\frac{1}{n} \sum_{i=1}^n \log \lambda_i\right) \quad (5)$$

where λ_i is the eigenvalue of Σ_{ij} .

C. Optimality Criteria On the Spectra

Recent research also investigates the optimality of the topological structure i.e. Laplacian matrix of the underlying pose graph. We start by representing the pose graph as a weighted weakly connected directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \omega)$,

where \mathcal{V} and \mathcal{E} are vertices and edges of the graph and ω are the weights for every edge in the graph. The connection relationship between every vertex and edge is defined by the incidence matrix $A = [a_1, a_2, \dots, a_m] \in \{-1, 0, 1\}^{n \times m}$. Then the Laplacian matrix corresponding to the directed graph \mathcal{G} is defined as:

$$L = A \text{Diag}\{\omega_1, \omega_2, \dots, \omega_m\} A^T. \quad (6)$$

The Laplacian matrix has many important properties, which can be derived from the graph connectivity from the analysis of the Laplacian spectrum. Consider that the eigenvalue of L is $\lambda_L^i = (\lambda_L^1, \lambda_L^2, \dots, \lambda_L^m)$, the second small eigenvalue in L , as known as the algebraic connectivity, is a key quantity metric as the inverse of the algebraic connectivity bounds the Cremér-Rao lower bound on the expected mean squared error for averaging [26]. Similarly, the optimality of the Laplacian matrix can also be defined using the number of spanning trees [21]. Given the number of spanning trees of \mathcal{G} :

$$S(\mathcal{G}) = \text{cof}(L) = \frac{1}{n} \prod_{i=2}^m \lambda_L^i, \quad (7)$$

where cof refers to the cofactor of L . Then the D-optimality of a Laplacian matrix is defined as:

$$D-opt(L) = (nS(\mathcal{G}))^{\frac{1}{n}}, \quad (8)$$

where n is the number of nodes.

IV. ACTIVE SLAM BASED ON SEMANTIC MUTUAL INFORMATION AND POSE GRAPH SPECTRUM

A. Semantic Mutual Information

In Section III-A, we have the methodology for computing mutual information within the occupancy grid map, which enables the robot to explore the trajectory that minimizes map uncertainty. While accounting for the robot kinematics and sensor errors on map entropy, equation (3) still falls short in active semantic mapping due to its limited capacity to handle semantic segmentation uncertainty. Deriving the representation of semantic mutual information is beneficial for improving the accuracy of the semantic map. For a map m with m_i where $m_i \in c, c \in [0, C]$ which means we have C different classes of segmentation (0 for free space), the PDF of m_i belongs to one specific class c is:

$$\begin{aligned} p(m_i = c) &= \text{logit}^{-1}(y) = \frac{e_{c+1}^T \exp(y_i)}{1^T \exp(y_i)} \\ e_{c+1}^T &= \begin{cases} 1 & e_{c+1}^{T[k]} = 1 \\ 0 & \text{other.} \end{cases} \end{aligned} \quad (9)$$

Here we use logit^{-1} to represent the inverse transform of log-odds, as the probability of y over c is expressed in a vector of log-odds:

$$y_i = \left[\log \frac{p(m_i = 0)}{p(m_i = 0)}, \dots, \log \frac{p(m_i = C)}{p(m_i = 0)} \right]. \quad (10)$$

The probability y is updated by:

$$y_{t+1,i} = y_{t,i} + \sum (\beta_i - y_{0,i}), \quad (11)$$

where β_i is the inverse observation model obtained by Bresenham's line algorithm [27]. To compute the semantic mutual information, instead of accumulating the MI in every grid which may be hit many times and considering that the ray from the sensor can intersect in multiple measurements, we follow the method used in [28] to select a subset of measurements in which there are no overlapping between rays, this leads to a low bound of the semantic mutual information. For a calculation time horizon T , the MI can be approximated as:

$$I(m; Z_{t:t+T, \Theta} | Z_{1:t-1, \Theta}) \geq \int_{\theta \in \Theta} I(m, \underline{Z}_{t:t+T, \theta}) d\theta \quad (12)$$

$$= \sum_{\tau=t}^{t+T} \sum_{\theta=1}^R \sum_{i \in I_\theta(l \max)} I(m_i; Z_{\tau, \theta} | Z_{1:t-1, \theta}).$$

where $\underline{Z}_{t:t+T, \theta}$ are the selected subset of measurements, R is the max indices of the ray, I_θ is the set of map cell indices along the ray and $l \max$ is the max sensing range. Using the expression in equation (3) and substituting (9) into (12), we have:

$$\int_{\theta \in \Theta} I(m, \underline{Z}_{t+1:t+T, \theta}) d\theta = \sum_{\tau=t}^{t+T} \sum_{\theta=1}^R \sum_{c=1}^C \int_0^{l \max} (p(Z_{t, \theta} = (l, c) | Z_{1:t-1}) \sum_{i \in I_\theta(l \max)} h(\beta_i(l, c)) - y_{0, i}, y_{t, i}) dl, \quad (13)$$

where

$$h(x, y) = \log\left(\frac{1^T \exp(y)}{1^T \exp(x + y)}\right) + y^T \text{logit}^{-1}(x + y). \quad (14)$$

To compute the first element in equation (13), let d denote the distance traveled by the ray within cell m_i , and i^* denote the index of the cell hit by $Z_{t, \theta}$, the first item is calculated by:

$$p(Z_{t, \theta} = (l, c) | Z_{1:t-1}) = \frac{p_t(m_{i^*} = c)}{d(i^*)} \prod_{i \in I_\theta(l) \setminus i^*} p_t(m_i = 0). \quad (15)$$

We utilize semantic mutual information as defined in (13) instead of relying on binary occupancy probabilities. This approach allows us to evaluate the confidence in the semantic classification of each voxel in the multi-class mapping. A voxel's category probability nearing 1 suggests a higher likelihood of belonging to a specific semantic category, which is also reflected in the reduction of uncertainty (entropy) in semantic maps.

B. Underlying Pose Graph Spectrum for Optimality Criteria

In addition to the dense mapping uncertainty, another important source of uncertainty comes from the robot's state estimation. It is even more important to analyze the error in pose graph optimization. According to the Cramér-Rao lower bound [26], the covariance matrix of any unbiased estimation

satisfies the lower bound defined as the Fisher Information Matrix (FIM). The FIM (also noted as Hessian matrix in some literature) is expressed as:

$$F(x) = \frac{1}{2} \sum_{k=1}^m J_k^T(x) \bar{\Sigma}^{-1} J_k(x), \quad (16)$$

where J is the Jacobian matrix of the error cost function with respect to the variable x and $\bar{\Sigma}^{-1}$ is the weighted covariance matrix. Recent research [21] has shown a close relationship between the FIM in the optimization and the Laplacian spectrum of the underlying pose graph structure. This relationship is linked with the different expression using the same FIM:

$$F(x) = \frac{1}{2} \sum_{k=1}^m I_k^T(x) \Sigma_k^{-1} I_k(x), \quad (17)$$

where

$$\Sigma_k^{-1} = Ad^T(T_k) \bar{\Sigma}_k^{-1} Ad(T_k), \quad (18)$$

where $Ad^T(T_k)$ is the adjoint matrix of the transformation T_k . We recommend readers to see [21] for a more detailed derivation. I_k is a sparse matrix consisting of several identity matrices \mathbb{I} filled in i -th rows and k -th columns. Based on this property, we can rewrite I_k in $I_k = a_k^T \otimes \mathbb{I}$ using the Laplacian generator a in section III-C and then equation (17) turns to:

$$F(x) = \frac{1}{2} \sum_{k=1}^m I_k^T(x) \Sigma_k^{-1} I_k(x) \propto \sum_{k=1}^m a_j a_j^T \otimes \Sigma_k^{-1} = \sum_{k=1}^m A_k \otimes \Sigma_k^{-1}. \quad (19)$$

The estimation uncertainty will keep increasing during exploration or decreasing when a loop closure optimization is performed, leading to a variable covariance matrix. Considering the upper-bound of Σ_k^{-1} , let $(\mu_1^k, \mu_2^k, \dots, \mu_n^k)$ the eigenvalue of Σ_k^{-1} in increasing order, it holds $\Sigma_k^{-1} \leq \mu_n^k \mathbb{I}$. Denote $\Phi_k = \Sigma_k^{-1}$, we will use the FIM Φ_k to weight every edge in the Laplacian graph. Using the infinite norm property and the concept of the Kronecker product \otimes , $\forall k, \Phi_k \leq \|\Phi_k\|_\infty$, the following inequality for (19) holds:

$$F \leq \sum_{k=1}^m (\|\Phi_k\|_\infty A_k) \otimes \mathbb{I} = L_\omega \otimes \mathbb{I}, \quad (20)$$

where L_ω is the weighted Laplacian graph with $\omega_k = \|\Phi_k\|_\infty$, using Weyl's monotonicity theorem [29], (20) turns to:

$$\|F\|_p \leq \|L_\omega\|_p, \quad \omega = \|\Phi_k\|_\infty, \forall p. \quad (21)$$

Equation (21) gives an upper-bound for $\|F\|_p$. Instead of using the infinite norm of Φ_k , we can relax the inequality relation by weighing every edge of the pose graph using the specific p -norm. We will get the following approximation:

$$\|F\|_p \approx \|L_\omega\|_p, \quad \omega = \|\Phi_k\|_p, \forall p. \quad (22)$$

Equation (22) gives a general relationship between the pose graph reliability and the underlying Laplacian spectrum, further, using (8), the optimality criteria is calculated by:

$$D\text{-opt}(F) \approx D\text{-opt}(L_\omega) = (nS(\mathcal{G}_\omega))^{\frac{1}{n}}. \quad (23)$$

C. Utility Function

Consistent with prior work [18], [19], we assume that the robot’s state and map uncertainties are independent, allowing for the utility function to be calculated as a heuristic linear weighting function. It is remarkable that in some recently advanced SLAM frameworks [30], [31] the uncertainty of the map points are also treated as a separate model and its influence on state estimation is also considered. Modeling map uncertainty during optimization differs from the active SLAM discussed in this paper. Following the classical exploration pipeline, we only take odometry results from SLAM estimation and we build the map for path planning in a separate dense mapping thread. We define the utility function in the form of Shannon-Rényi entropy as in literature [18], however, we further improve it in our approach by evaluating all mutual information along each potential path, instead of calculating the proportion of unknown cells only at the frontier. This approach helps to mitigate the stochasticity associated with frontier detection and incorporates the effect of semantic uncertainty and also promotes the selection of trajectories that can improve the overall semantic mapping accuracy, not just at the frontier. The utility function is defined as:

$$\begin{aligned} a^* &= \arg \max_a SR(a, T) \\ &= \arg \max(D\text{-opt}(L_\omega^{a,T}) - (\frac{1}{1-\alpha})D\text{-opt}(L_\omega^{a,T})), \end{aligned} \quad (24)$$

where

$$\alpha = 1 + \frac{Cost(a, T)}{\int_{\theta \in \Theta} I(m, \underline{Z}_{t+1:t+T}, \theta) d\theta} \quad (25)$$

where a^* is the best action among all potential actions a , SR refers to Shannon-rényi entropy, T is the action time, $(D\text{-opt}(L_\omega^{a,T})$ and $\int_{\theta \in \Theta} I(m, \underline{Z}_{t+1:t+T}, \theta) d\theta$ can be calculated by (23) and (13), respectively. Parameter α essentially means using expected information gain to discount the localization error from a trajectory. $Cost(a, T)$ refers to the cost for a path, in this paper, we use traveling distance. Thanks to the properties of the Shannon-Rényi entropy, (24) circumvents the drawbacks associated with manually adjusting the weighting parameters in a simple linear weighting function. This approach effectively balances the exploration of new regions—aiming to achieve higher information gains and reduce map uncertainty—with the strategic revisiting of previously explored areas, based on robot localization uncertainty.

Notably, the potential for loop closure allows the robot to reduce its localization uncertainty without revisiting previously explored areas. We take this into account when calculating the D-optimality. Specifically, while weighting the underlying Laplacian graph, those edges likely to form

loop closure are assigned greater weight rather than only calculating the FIM. This results in higher D-optimal values for such edges, which implicitly encourages the robot to actively search for loop closures.

D. Autonomous Exploration

To perform autonomous exploration, we first project the 3D map onto a 2D space. Subsequently, we delineate the boundary between explored and unexplored regions, utilizing edge detection. We then cluster map cells that share connected boundary components, with each cluster representing a frontier. For each frontier, we use the A^* algorithm to generate a path to each center of the frontier. All generated paths will be evaluated using the proposed function. Goals that cannot be reached, either due to the impossibility of path computation or because they were previously identified as frontiers, are blacklisted, indicating that the goal is currently unreachable. From the global path determined by A^* , we select several nodes as waypoints for the local planner, lattice space search [32] is used for local planning and obstacle avoidance. Finally, the underlying controller generates the control commands based on the output of the local planner.

V. EXPERIMENTS

A. Environment Setup

We evaluate our method in a photorealistic simulation environment [33]. The simulation experiments are conducted on a laptop with an Intel Core i5-11400h CPU and NVIDIA GeForce RTX 3060 GPU. The robot is equipped with an RGB-D camera and the semantic segmentation comes from the Habitat simulator [34]. The semantic segmentation is running at 640×360 resolution with a camera frame rate of 30hz. We adopt the graph-based visual SLAM approach in [35], leveraging the g2o library [36] for backend optimization. We use a local planner in [33] with 3D terrain analysis capabilities. Our environment representation is based on the semantic octomap [37], within which we evaluate semantic mutual information in a 3D space. The map resolution is selected to 0.25 and the semantic mapping is running at 2hz. Although the Habitat simulator is primarily designed for 2D tasks, our approach can be easily extended to 3D environments.

We conducted 10 repetitions of the exploration in two different environments on the Matterport dataset [38], with exploration stopping criteria of 5 minutes for the first smaller building and 10 minutes for the second larger building. We compared our approach with the state-of-the-art methods SSMI [2] and TARE [12], where SSMI is an advanced active semantic mapping method based on information theory, and TARE has been applied in large-scale outdoor environments like the DARPA SubT Challenge. We used the default parameters in their settings. We focus on the localization accuracy and the quality of the semantic maps during exploration, they will be analyzed in section V-B. In addition, we also compare the exploration efficiency of each method by monitoring the explored area volume over time.

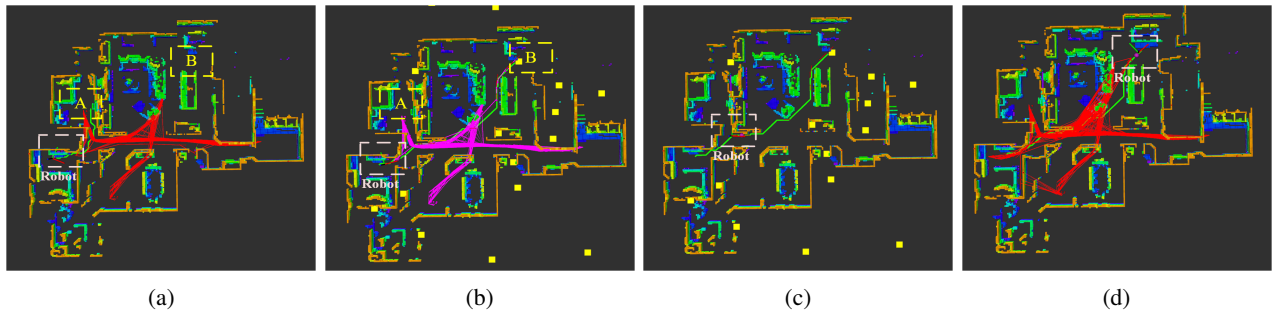


Fig. 2: Schematic diagram of a decision-making process during the active SLAM. The path of frontier A has shorter travel distance while the path of frontier B will provide a more optimistic pose graph with more co-visibility points and loop closure edges for visual SLAM. The final choice of exploration is frontier B. In the figure, red and pink lines represent pose graphs at different stages. Yellow rectangles are candidate frontiers, some of them have been excluded before the decision due to inaccessibility. Green line represents the global path planned by A* algorithm and pink rectangle represents the robot. (a) The robot plans a path toward frontier A and evaluates the corresponding hallucination graph and semantic mutual information. (b) The robot plans a path toward frontier B and evaluates the corresponding hallucination graph and semantic mutual information. (c) The robot executes a path to frontier B. (d) The final generated pose graph after the robot reaches frontier B.

B. Results and Discussion

1) *Geometric SLAM Metric Evaluation:* We employed the widely used Absolute Trajectory Error (ATE) as the localization evaluation metric and the mean Euclidean distance as the metric for map evaluation as in [31].

In practice, we recorded the data from each exploration in the rosbag format. Localization errors are computed using evo [39]. For semantic mapping, we saved the maps in binary color octomap format, converted into point cloud maps at a resolution of 0.25m for evaluation using the CloudCompare software [40]. We plot the distribution of the localization and map errors for each experiment as shown in Fig. 3. Compared to the other two methods, our method exhibits a reduction in average localization and map errors with a maximum of 38% and 21%. Specifically, our method obtained the smallest localization and map errors in 8 of 10 repeated times in experiment 1, and 9 of 10 in experiment 2. This is attributed to our approach’s focus on uncertainty in SLAM pose graph optimization, enabling the robot to intentionally (re)select trajectories with less uncertainty instead of greedily traveling to the place with less traveling distance or information gain. See Fig. 2.

2) *Semantic Mapping and Classification Results:* As the evaluation method in [41], we use Intersection Over Union (IoU) as the semantic map evaluation metric. It is worth noting that the simulator itself does not provide any semantic map, so we initially built a ground truth 3D semantic map based on noise-free state estimation and semantic segmentation from the simulator. Subsequently, we compare the ground truth map with the semantic map constructed upon completion of exploration, as shown in Fig. 4. The Matterport dataset has 40 instance labels, we focus on the top eight most frequently appearing categories for comparison. The quantitative results are given in Tables I and II.

Our method achieves the highest average IoU values across these semantic categories, outperforming other base-

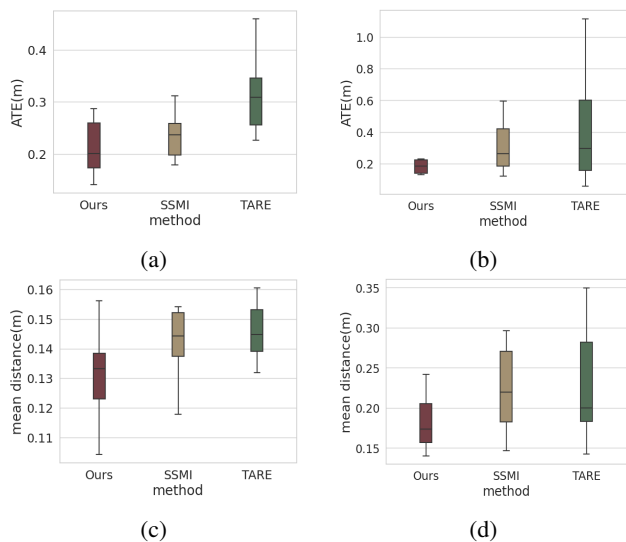


Fig. 3: SLAM error results for 10 repetitions of the exploration experiment. (a) Localization error in experiment 1. (b) Localization error in experiment 2. (c) Map error in experiment 1. (d) Map error in experiment 2.

TABLE I: Quantitative results of experiment 1 for 8 common semantic classes.

Metric	Method	objects	table	cabinet	wall	clothes	mirror	picture	lighting	Average
IoU	Ours	47.46	47.27	46.21	45.12	40.87	38.15	37.56	34.56	42.15
	SSMI	45.25	48.01	45.25	42.71	37.55	38.08	36.47	32.42	40.72
	TARE	39.62	45.47	39.74	25.35	37.82	23.04	30.34	21.18	32.82

TABLE II: Quantitative results of experiment 2 for 8 common semantic classes.

Metric	Method	ceiling	appliances	sink	stool	plant	counter	table	clothes	Average
IoU	Ours	50.73	45.26	43.91	43.51	40.42	39.18	37.99	33.57	41.82
	SSMI	46.02	41.01	25.13	43.80	39.30	36.12	29.25	28.83	36.18
	TARE	42.01	36.86	23.86	41.17	32.51	31.70	23.27	30.67	32.76




Fig. 4: Visualization of metric-semantic map results obtained during one exploration experiment. (a) Partially explored metric map. (b) Ground truth metric map. (c) The error (Euclidean distance) in CloudCompare software calculation. (d) Ground truth map visualization for the environment. (e) Partially explored semantic map. (f) Ground truth semantic map.

line methods with up to a 9.33% improvement. Our method obtained the highest average IoU value 8 of 10 repeated times in both experiments 1 and 2. The average IoU values of SSMI are higher than those of TARE, as SSMI also takes semantic uncertainty into account. Another interesting observation is that our IoU values match, but are slightly lower than the overall level of multi-class mapping results reported in [41]. This discrepancy can be attributed to our utilization of a naive discrete voxel probability fusion updating approach. Integrating the active mapping method proposed in our paper with other, more advanced semantic mapping methods, such as [41], could offer further improvements.

3) *Exploration Metrics*: Finally, we investigated the impact of evaluating semantic mutual information as well as pose graph optimality on exploration efficiency. This involved plotting the volume explored over time across multiple iterations of our experiments within two different environments, as illustrated in Fig. 5. In the plots, the middle line denotes the median value, the upper and lower bounds are represented by the shadow area. Interestingly, we observed partial loss of localization results during some experiments conducted in the baselines, for a fair comparison we excluded these experiments and focused only on those where the entire exploration stack worked well. The first conclusion from the results is that our method closely matches SSMI in terms of exploration efficiency, albeit with a slight increase in time overhead. This trade-off is considered advantageous, as it

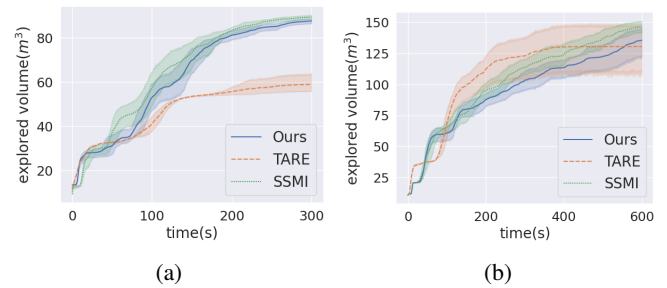


Fig. 5: Comparison of exploration results. (a) The curves of explored volume over time in experiment 1. (b) The curves of explored volume over time in experiment 2.

aims to enhance SLAM performance and reduce the risk of localization failures, such as tracking lost, which could significantly impact the exploration system.

The second observation is that the information theory-based exploration strategies (our approach and SSMI) outperform TARE in small-scale environments. However, TARE shows better performance in larger environments due to its hierarchical exploration strategy, which is specifically designed to address large-scale exploration challenges. This work aims to provide a mathematical explanation and complement the state-of-the-art exploration methods from an uncertainty perspective, rather than offering a replacement. Using a combination of global and local exploration strategies based on different resolutions rather than information-theoretic methods would be beneficial in accelerating exploration, although this is outside the scope of this paper.

Last, it is also important to mention that as the pose graph grows, the overhead of computing the underlying weighted Laplacian graph may be significant. At this point, spectral measurement sparsification for the pose graph may be a promising tool to further improve exploration efficiency.

VI. CONCLUSION

In this paper, we proposed a real-time active metric-semantic SLAM method for efficient and accurate robot exploration. Our method not only evaluates the path with the largest semantic information gain in 3D space, but also integrates the effect of robot state uncertainty. The computation of the large and dense FIM is transformed into the analysis of the underlying pose graph topology, allowing planning to be executed online. The Shannon-Rényi formed entropy avoids different orders of magnitude of uncertainty, leading to an autonomous balance between exploration and exploitation. Comparison with advanced exploration methods including SSMI and TARE demonstrates the performance of our implementation. The research on semantic uncertainty is also meaningful for computer vision tasks, such as mitigating the impact of domain shifts caused by data and scenes. Our implementation can also be used for aerial robot exploration, with a 3D path planner. Future work will focus on spectral graph sparsification to accelerate exploration. We are also interested in generalizing to multi-robot exploration and conducting field experiments.

REFERENCES

- [1] Z. Zhang, T. Henderson, S. Karaman, and V. Sze, "Fsmi: Fast computation of shannon mutual information for information-theoretic mapping," *The International Journal of Robotics Research*, vol. 39, no. 9, pp. 1155–1177, 2020. [Online]. Available: <https://doi.org/10.1177/0278364920921941>
- [2] A. Asgharivaskasi and N. Atanasov, "Semantic octree mapping and shannon mutual information computation for robot exploration," *IEEE Transactions on Robotics*, 2023.
- [3] J. McCormac, A. Handa, A. Davison, and S. Leutenegger, "Semanticfusion: Dense 3d semantic mapping with convolutional neural networks," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*, 2017, pp. 4628–4635.
- [4] M. Grinvald, F. Furrer, T. Novkovic, J. J. Chung, C. Cadena, R. Siegwart, and J. Nieto, "Volumetric Instance-Aware Semantic Mapping and 3D Object Discovery," *IEEE Robotics and Automation Letters*, vol. 4, no. 3, pp. 3037–3044, July 2019.
- [5] A. Rosinol, M. Abate, Y. Chang, and L. Carlone, "Kimera: an open-source library for real-time metric-semantic localization and mapping," in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2020. [Online]. Available: <https://github.com/MIT-SPARK/Kimera>
- [6] A. Rosinol, A. Gupta, M. Abate, J. Shi, and L. Carlone, "3D dynamic scene graphs: Actionable spatial perception with places, objects, and humans," in *Robotics: Science and Systems (RSS)*, 2020.
- [7] Q. Gu, A. Kuwajerwala, S. Morin, K. Jatavallabhula, B. Sen, A. Agarwal, C. Rivera, W. Paul, K. Ellis, R. Chellappa, C. Gan, C. de Melo, J. Tenenbaum, A. Torralba, F. Shkurti, and L. Paull, "Conceptgraphs: Open-vocabulary 3d scene graphs for perception and planning," *arXiv*, 2023.
- [8] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark *et al.*, "Learning transferable visual models from natural language supervision," in *International conference on machine learning*. PMLR, 2021, pp. 8748–8763.
- [9] A. Rosinol, J. J. Leonard, and L. Carlone, "Nerf-slam: Real-time dense monocular slam with neural radiance fields," in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2023, pp. 3437–3444.
- [10] N. Keetha, J. Karhade, K. M. Jatavallabhula, G. Yang, S. Scherer, D. Ramanan, and J. Luiten, "SplaTAM: Splat, track & map 3D gaussians for dense RGB-D SLAM," *arXiv preprint*, 2023.
- [11] T. Dang, M. Tranzatto, S. Khattak, F. Mascari, K. Alexis, and M. Hutter, "Graph-based subterranean exploration path planning using aerial and legged robots," *Journal of Field Robotics*, vol. 37, no. 8, pp. 1363–1388, 2020, Wiley Online Library.
- [12] C. Cao, H. Zhu, Z. Ren, H. Choset, and J. Zhang, "Representation granularity enables time-efficient autonomous exploration in large, complex worlds," *Science Robotics*, vol. 8, no. 80, p. eadf0970, 2023.
- [13] B. Charrow, S. Liu, V. Kumar, and N. Michael, "Information-theoretic mapping using cauchy-schwarz quadratic mutual information," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*, 2015, pp. 4791–4798.
- [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," 2023.
- [15] 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*, 2023.
- [16] L. Carlone, J. Du, M. Kaouk Ng, B. Bona, and M. Indri, "Active slam and exploration with particle filters using kullback-leibler divergence," *Journal of Intelligent & Robotic Systems*, vol. 75, pp. 291–311, 2014.
- [17] H. Carrillo, I. Reid, and J. A. Castellanos, "On the comparison of uncertainty criteria for active slam," in *2012 IEEE International Conference on Robotics and Automation*, 2012, pp. 2080–2087.
- [18] H. Carrillo, P. Dames, V. Kumar, and J. A. Castellanos, "Autonomous robotic exploration using occupancy grid maps and graph slam based on shannon and rényi entropy," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*, 2015, pp. 487–494.
- [19] J. A. Placed and J. A. Castellanos, "Fast autonomous robotic exploration using the underlying graph structure," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 6672–6679.
- [20] Y. Tao, X. Liu, I. Spasojevic, S. Agarwal, and V. Kumar, "3d active metric-semantic slam," *IEEE Robotics and Automation Letters*, 2024.
- [21] J. A. Placed and J. A. Castellanos, "A general relationship between optimality criteria and connectivity indices for active graph-slam," *IEEE Robotics and Automation Letters*, vol. 8, no. 2, pp. 816–823, 2022.
- [22] J. A. Placed, J. J. G. Rodríguez, J. D. Tardós, and J. A. Castellanos, "Explorb-slam: Active visual slam exploiting the pose-graph topology," in *Iberian Robotics conference*. Springer, 2022, pp. 199–210.
- [23] Y. Tian and J. P. How, "Spectral sparsification for communication-efficient collaborative rotation and translation estimation," *IEEE Transactions on Robotics*, 2023.
- [24] K. J. Doherty, D. M. Rosen, and J. J. Leonard, "Spectral measurement sparsification for pose-graph slam," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 01–08.
- [25] P.-Y. Lajoie and G. Beltrame, "Swarm-slam: Sparse decentralized collaborative simultaneous localization and mapping framework for multi-robot systems," *IEEE Robotics and Automation Letters*, vol. 9, no. 1, pp. 475–482, 2023.
- [26] Y. Chen, S. Huang, L. Zhao, and G. Dissanayake, "Cramér-rao bounds and optimal design metrics for pose-graph slam," *IEEE Transactions on Robotics*, vol. 37, no. 2, pp. 627–641, 2021.
- [27] J. E. Bresenham, "Algorithm for computer control of a digital plotter," *IBM Systems Journal*, vol. 4, no. 1, pp. 25–30, 1965.
- [28] A. Asgharivaskasi and N. Atanasov, "Active bayesian multi-class mapping from range and semantic segmentation observations," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 1–7.
- [29] I. Gohberg and M. G. Kren, *Introduction to the theory of linear nonselfadjoint operators*. American Mathematical Soc., 1978, vol. 18.
- [30] Z. Chen, Y. Xu, S. Yuan, and L. Xie, "ig-lio: An incremental gicp-based tightly-coupled lidar-inertial odometry," *IEEE Robotics and Automation Letters*, vol. 9, no. 2, pp. 1883–1890, 2024.
- [31] X. Hu, L. Zheng, J. Wu, R. Geng, H. Wei, L. Wang, J. Jiao, and M. Liu, "Paloc: Advancing slam benchmarking with prior-assisted 6-dof trajectory generation and uncertainty estimation," *IEEE/ASME Transactions on Mechatronics*, 01 2024.
- [32] J. Zhang, C. Hu, R. G. Chadha, and S. Singh, "Falco: Fast likelihood-based collision avoidance with extension to human-guided navigation," *Journal of Field Robotics*, vol. 37, no. 8, pp. 1300–1313, 2020.
- [33] [Online]. Available: <https://www.cmu-exploration.com/development-environment>
- [34] A. Szot, A. Clegg, E. Undersander, E. Wijmans, Y. Zhao, J. Turner, N. Maestre, M. Mukadam, D. Chaplot, O. Maksymets, A. Gokaslan, V. Vondrus, S. Dharur, F. Meier, W. Galuba, A. Chang, Z. Kira, V. Koltun, J. Malik, M. Savva, and D. Batra, "Habitat 2.0: Training home assistants to rearrange their habitat," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- [35] R. Mur-Artal and J. D. Tardós, "Orb-slam2: An open-source slam system for monocular, stereo, and rgb-d cameras," *IEEE Transactions on Robotics*, vol. 33, no. 5, pp. 1255–1262, 2017.
- [36] R. Kümmerle, G. Grisetti, H. Strasdat, K. Konolige, and W. Burgard, "g 2 o: A general framework for graph optimization," in *2011 IEEE international conference on robotics and automation*. IEEE, 2011, pp. 3607–3613.
- [37] A. Hornung, K. M. Wurm, M. Bennewitz, C. Stachniss, and W. Burgard, "OctoMap: An efficient probabilistic 3D mapping framework based on octrees," *Autonomous Robots*, 2013, software available at <https://octomap.github.io>. [Online]. Available: <https://octomap.github.io>
- [38] A. Chang, A. Dai, T. Funkhouser, M. Halber, M. Niessner, M. Savva, S. Song, A. Zeng, and Y. Zhang, "Matterport3d: Learning from rgb-d data in indoor environments," *International Conference on 3D Vision (3DV)*, 2017.
- [39] M. Grupp, "evo: Python package for the evaluation of odometry and slam." <https://github.com/MichaelGrupp/evo>, 2017.
- [40] [Online]. Available: <https://www.danielgm.net/cc/t>
- [41] L. Gan, R. Zhang, J. W. Grizzle, R. M. Eustice, and M. Ghaffari, "Bayesian spatial kernel smoothing for scalable dense semantic mapping," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 790–797, April 2020.