

Enough is Enough: Towards Autonomous Uncertainty-driven Stopping Criteria ^{*}

Julio A. Placed ^{*} José A. Castellanos ^{*}

^{*} *Departamento de Informática e Ingeniería de Sistemas, Instituto de Investigación en Ingeniería de Aragón, Universidad de Zaragoza, C/ María de Luna 1, 50018, Zaragoza, Spain. (e-mail: {jplaced, jacaste}@unizar.es)*

Abstract: Autonomous robotic exploration has long attracted the attention of the robotics community and is a topic of high relevance. Deploying such systems in the real world, however, is still far from being a reality. In part, it can be attributed to the fact that most research is directed towards improving existing algorithms and testing novel formulations in simulation environments rather than addressing practical issues of real-world scenarios. This is the case of the fundamental problem of autonomously deciding when exploration has to be terminated or changed (*stopping criteria*), which has not received any attention recently. In this paper, we discuss the importance of using appropriate *stopping criteria* and analyse the behaviour of a novel criterion based on the evolution of *optimality criteria* in active graph-SLAM.

Copyright © 2022 The Authors. This is an open access article under the CC BY-NC-ND license (<https://creativecommons.org/licenses/by-nc-nd/4.0/>)

Keywords: Active SLAM, autonomous exploration, stopping criteria, decision making, graph-SLAM

1. INTRODUCTION

Actively exploring unknown environments has attracted the attention of the robotics community since the very first works appeared (Bajcsy, 1988; Whaite and Ferrie, 1997). Basically, it encompasses controlling the movement of a robot which is capable of building a representation of the environment and localising itself on it. The latter are known as Simultaneous Localisation and Mapping (SLAM) and is the foundation of robotic exploration. See Thrun et al. (2005), Durrant-Whyte and Bailey (2006), Cadena et al. (2016) and references there in. Specifically, this article is concerned with graph-based SLAM, in which the problem is formulated using a graphical representation that encodes the robot poses in the graph nodes and the constraints between them (e.g. sensor measurements) in edges (Grisetti et al., 2010).

Active SLAM (also referred to as robotic exploration) augments the above problem by also including a planning step. Thus, the autonomous agents must be capable of deciding where they should move next in order to improve the model of the environment (i.e. the map built and the localisation accuracy). Typically, this decision making problem has been viewed from an uncertainty-reduction standpoint, and divided in three phases for tractability: identification of possible actions, evaluation of their usefulness and execution of those considered as optimal. Solving this paradigm in finite time usually requires making several assumptions (e.g. maximum likelihood observations (Platt et al., 2010; Valencia and Andrade-Cetto, 2018; Placed and Castellanos, 2021b) or a finite action

set (Carrillo et al., 2018; Placed and Castellanos, 2020) but, even so, the second stage is computationally time-consuming: comparing the utility of two (sets of) actions involves reasoning over the probability density functions of the two random variables of interest, i.e. over the (usually high-dimensional) covariance or Fisher information matrices (FIM) of the beliefs.

In this sense, identifying the moment in which the exploration strategy is no longer adding information to the system is of utmost importance, in order to avoid unnecessary computational load. At that point, it could be changed or the exploration could simply be terminated. Besides obtaining irrelevant information, the continuous acquisition of redundant data has demonstrated to be detrimental for SLAM and may lead to map inconsistencies and even unrecoverable states. On the other hand, the ability to determine task completion is key for robots in real environments to be autonomous. Constraints typical of simulation (e.g. percent of coverage) are no longer valid in real-world applications in which the environment to explore is unknown by definition. The set of rules for deciding whether or not to terminate exploration is known as *stopping criteria* (SC) or *termination criteria*, and it was already identified by Cadena et al. (2016) as an open challenge more than five years ago. Despite the great impact of this work, however, no further research has been done in this field.

In this paper, we stress the importance of using task-driven SC towards autonomous robotic exploration, and thoroughly discuss why those commonly used in the literature are neither meaningful nor practical for that purpose. In addition, following the idea of using evolution of *optimality criteria* as SC (Cadena et al., 2016) and their fast computation via graph connectivity indices (Placed

^{*} This work was supported by the Spanish government under grant PID2019-108398GB-I00 and by Aragón government under grant DGA_FSE T45_20R.

and Castellanos, 2021b), we present and analyse the performance of a novel criterion based on D -optimality. We seek to stimulate this field of research and encourage the use of meaningful criteria in future active SLAM works.

The rest of the manuscript is organised as follows. Section 2 presents the problem tackled. Section 3 contains a discussion on the different *stopping criteria* used in the literature and their limitations. In section 4, we describe and validate the proposed approach. Finally, the paper is concluded in Section 5, also outlining future work.

2. BACKGROUND

2.1 Active graph-SLAM

Active SLAM refers to the problem of controlling the movement of a robot which is performing SLAM in an unknown environment with the main objective of improving the quality of the SLAM estimates (\mathcal{S}): the map representation and the robot state. This paradigm requires agents to be capable of autonomously selecting the optimal actions to execute. Under the wider framework of Partially Observable Markov Decision Processes (POMDP), it can be formulated as a decision making problem in which a decision, δ , is defined by the mapping:

$$\delta : \mathcal{P}(\mathcal{S}) \rightarrow \mathcal{A} \quad (1)$$

That is, it maps elements from the space of probabilities over the robot and map states (i.e. beliefs, $\mathcal{P}(\mathcal{S})$) to the action space, \mathcal{A} .

In general, a decision will be preferred over another if it is more useful. Utility functions (\mathcal{U}) measure this usefulness by mapping the complex spaces $\mathcal{P}(\mathcal{S})$ and \mathcal{A} to the real scalar space, in which comparison can be made directly:

$$\mathcal{U} : \mathcal{P}(\mathcal{S}) \times \mathcal{A} \rightarrow \mathbb{R} \quad (2)$$

Thus, decision making in active SLAM can be reduced to the following optimisation problem:

$$\mathbf{a}^* = \arg \max_{\mathbf{a} \in \mathcal{A}} \mathcal{U} \quad (3)$$

where \mathbf{a}^* is the optimal set of actions to be executed.

Traditionally, active SLAM has been divided in three differentiated phases for the ease of its resolution (Makarenko et al., 2002), which are directly linked to eq. (1)-(3). During first step, and just for computational purposes, the (usually finite) set of possible destinations the robot could travel to is identified (e.g. via frontier detection). Second stage deals with utility computation. The effect of executing the candidate actions is evaluated by quantifying the expected uncertainty in the two random variables. Information-theoretic (IT) approaches are the most common and are usually based on the concept of entropy (Stachniss and Burgard, 2003; Valencia et al., 2012; Palomeras et al., 2019). On the other hand, methods based on Theory of Optimal Experimental Design (TOED) directly compute utility from the variance of the variables of interest. *Optimality criteria* quantify utility by mapping elements from the space of real symmetric matrices (covariance or FIM) to \mathbb{R} . D -optimality has been considered the most fruitful criterion (Rodríguez-Arévalo et al., 2018), and its modern formulation is as follows (Pukelsheim, 2006):

$$D\text{-opt} \triangleq \exp \left(\frac{1}{d} \sum_{k=1}^d \log(\lambda_k) \right) \quad (4)$$

where $\Lambda = (\lambda_1, \dots, \lambda_d)$ are the eigenvalues of the matrix under study, with dimension d . Finally, during the third step, the frontier with highest utility is selected as optimal and the set of actions to reach it is executed. The entire sequence would be repeated until a certain SC was met.

This paper is particularly focused on active graph-SLAM methods, in which SLAM is formulated using a graph representation where nodes represent the robot poses, and edges encode the constraints between them, usually in the form of a FIM. For the j -th edge, it will be $\Phi_j \in \mathbb{R}^{\ell \times \ell}$ with ℓ the dimension of the state vector (e.g. 3 for 2D SLAM). These methods rely on the insight that the map representation (usually in the form of an occupancy grid) can be retrieved once the robot states have been properly estimated (Montemerlo and Thrun, 2003; Grisetti et al., 2010). Evaluation of *optimality criteria* in such cases involves computing the eigenvalues of the system's FIM, $\mathbf{Y} \in \mathbb{R}^{n\ell \times n\ell}$, with n the number of nodes in the graph. It should be noted that the matrix dimensions grow as the trajectory does, and that the map information must somehow be included in \mathbf{Y} to effectively address the task under study. See Placed and Castellanos (2021b) for a more detailed discussion on active graph-SLAM and \mathbf{Y} .

2.2 Exploiting the Graphical Structure in Active SLAM

Evaluation of utility is costly and quickly becomes intractable for online approaches. It usually represents the main bottleneck in active SLAM, especially if TOED-based. Thus, most works resort to IT, sparse information matrices (Indelman et al., 2015) or sparsified representations (Carrillo et al., 2018; Elimelech and Indelman, 2019).

Regarding computational complexity, the prominent works by Khosoussi et al. (2014, 2019) show that classical *optimality criteria* are closely linked to the connectivity of the underlying pose-graph. Thus, instead of evaluating the whole FIM, decision making in active graph-SLAM can be done in a fraction of the time—up to an order of magnitude faster (Kitanov and Indelman, 2019; Placed and Castellanos, 2021b)—by studying the graph connectivity indices. Placed and Castellanos (2021a,b) formulate the existing relationship between modern *optimality criteria* of the FIM and those of the weighted Laplacian, for the general case of active graph-SLAM. Specifically, for D -optimality:

$$D\text{-opt}(\mathbf{Y}) \approx (n t(\mathcal{G}_\gamma))^{\frac{1}{n}} \quad (5)$$

where \mathcal{G}_γ denotes a pose-graph in which each edge is weighted with the same criterion to be estimated, i.e. $\gamma_j = D\text{-opt}(\Phi_j)$, and $t(\mathcal{G}_\gamma)$ is the weighted number of spanning trees of the graph, equal to any cofactor of its weighted Laplacian. See Placed and Castellanos (2021b).

3. LIMITATIONS OF THE EXISTING STOPPING CRITERIA

Most active SLAM works in the literature resort to spatial or temporal constraints as SC. The former can be used only because they use known (simulation) environments

Table 1. Example works, formulation basis and limitations of the most relevant SC in literature.

Stopping Criterion		Works	Formulation	Limitations
Temporal		Leung et al. (2008), Valencia et al. (2012), Carrillo et al. (2018) and Placed and Castellanos (2021b)	Time space	<ul style="list-style-type: none"> • No task completion • Must be set experimentally (scenario- and robot-dependent) • Not comparable across systems
Spatial	Frontiers	Yamauchi et al. (1999), Korb and Schöttl (2018) and Pimentel et al. (2018)	State space	<ul style="list-style-type: none"> • No task completion (coverage only) • Difficult to compare across systems • Unobservable and unreachable areas • Most times require prior knowledge of the environment or human interaction
	Coverage	Pham and Juang (2013), Amigoni et al. (2013), Lenac et al. (2016), Palomeras et al. (2019), Chen et al. (2020) and Xu et al. (2022)		
	Other	Bircher et al. (2018) and Suresh et al. (2020)		
Uncertainty-based	IT	Simmons et al. (2000), Stachniss et al. (2004), Salan et al. (2014), Chen et al. (2019), Ghaffari Jadidi et al. (2019), Gomez et al. (2019) and Deng et al. (2020)	Information space	<ul style="list-style-type: none"> • No task completion (coverage only) • Difficult to compare across systems • Unobservable and unreachable areas • May require prior knowledge of the environment
	TOED	—	Task space	<ul style="list-style-type: none"> • Computational and formulation complexity

for which the complete size of the map is available, and the latter typically encode physical constraints of the robot (e.g. battery) rather than those of the variables of interest. Both of them must be experimentally designed for a specific scenario and can be used neither in unknown environments, nor across systems. Table 1 illustrates the preferred choice of SC in some relevant works. It is easy to notice how temporal and spatial (i.e. geometric) criteria monopolise the attention in the literature.

Limiting the amount of time for an exploration task has been used extensively (Leung et al., 2008; Valencia et al., 2012; Carrillo et al., 2018; Placed and Castellanos, 2021b) for two main reasons: its simplicity and because comparing the performance of diverse agents in the same time horizon is straightforward. This criterion does not directly require previous knowledge about the environment, although for exploration results to be relevant, they have to be accordant. In any case, it (i) does not guarantee exploration of the entire environment or low uncertainty estimates (i.e. task-completion), (ii) is dependent on the scenario and the robot configuration, and (iii) often needs to be set experimentally (e.g. by manually exploring the environment before testing the approach).

The simplest geometric SC broadly used in the literature is the non-existence of locations to explore, which usually translates in the absence of (clusters of) frontiers (Yamauchi et al., 1999; Korb and Schöttl, 2018; Pimentel et al., 2018). A second relevant geometric constraint is the desired coverage (e.g. 90%). For its use, by definition, it is mandatory to know the size or even the full map of the environment. Despite being in conflict with active SLAM definition, it is often used (Pham and Juang, 2013; Amigoni et al., 2013; Lenac et al., 2016; Palomeras et al., 2019; Chen et al., 2020). A recent example can be found in (Xu et al., 2022), where they propose a unified framework to benchmark robotic exploration, although assume known environments and base efficiency on coverage. Other geometric-related SC rely on limiting the size of the exploration trees (Bircher et al., 2018), or on qualitative

metrics that require human supervision (Suresh et al., 2020). In contrast to temporal SC, all of the above allow checking the completion of the task of covering a surface or volume, since they are formulated over the state space. Nevertheless, two insights are worth noting: (i) the fundamental aspect in active SLAM of lack of prior information about the environment is violated, and (ii) they do not assess the active SLAM task but a coverage one, so they have no concern about the quality of the estimates. In fact, they do not even monitor the robot localisation.

Following Information Theory (IT) and the well-known (approximation) relationship between the entropy of an occupancy map and the number of unknown cells on it (Stachniss, 2009), a number of works (Simmons et al., 2000; Chen et al., 2019; Deng et al., 2020) soon reformulated the aforementioned spatial SC in the so-called information space (or belief space). That is, the place where the probability density functions of the random variables naturally live. See Barraquand and Ferbach (1995) and Prentice and Roy (2009). These probability distributions reflect knowledge about the system, and related metrics typically measure the entropy reduction in the posterior (map) distribution. Gomez et al. (2019) offer a fresher look at spatial criteria, using the concept of “interesting” frontiers, that is, frontiers with an expected utility above a defined threshold (e.g. expected Information Gain, IG). Determining this threshold must be done experimentally, though, and is scenario-specific. Ghaffari Jadidi et al. (2019) use a saturation information value over the current (and not necessarily complete) map, which reflects the confidence in the representation. Stachniss et al. (2004) observe that considering only a constant upper bound on the map information can lead to repeatedly acquiring the same information, and that is indeed detrimental to the SLAM performance. Salan et al. (2014) also realise that coverage SC are impractical, even when formulated over the information space: some cells can be unobservable or even unreachable. Instead, they propose to conduct exploration until there are no possible configurations that maximised information.

It is essential to note that, despite fast, information-driven metrics do not allow to check for task-completion in any case, since they are not related to the task of active SLAM but to that of coverage. Consider the example case in which all potential locations to travel to would result in low informative returns (e.g, entropy or IG below a predefined threshold or even zero). According to information-driven SC, exploration would be finished, but not necessarily implying that the two variables of interest were perfectly nor even well estimated, since the information units are unrelated to the physical meaning of the task (e.g. m^2 for linear variance).

4. TOWARDS MEANINGFUL STOPPING CRITERIA

The use of metrics (either or raw or their evolution) that stem from TOED as *stopping criteria* has been identified as promising many times (Cadena et al., 2016; Lluvia et al., 2021), although no research has been done in the topic. In contrast to IT metrics, they are task-driven, that is, they evaluate uncertainty in the task space (i.e. variance of estimates) and therefore allow to directly check if a given set of actions improves the task or if the estimations are sufficiently accurate. This also implies that they can be compared across systems. Their main drawback is the high computational load that evaluating large covariance matrices has traditionally required. However, thanks to the fast computation of graph connectivity indices and their equivalence to *optimality criteria* in active graph-SLAM (see Section 2.2), we propose their evolution along exploration as SC.

A well-conditioned SC has to account for both variables of interest in active SLAM: the robot localisation and map representation. Revisiting Section 2.1, graph-based SLAM builds upon the idea that the map can be straightforwardly retrieved once the robot is properly localised. Therefore, we propose the following SC based on the evolution of the robot's uncertainty (via graph's D -optimality) and the mapped area:

$$\Gamma = \Delta\mathcal{U} + |\Delta A| < \Gamma_{th} \quad (6)$$

where \mathcal{U} is the utility function (see equation 5), A the known area in the map and Δ denote percentage variations w.r.t. previous active SLAM steps. The area has been restricted to absolute variations, as otherwise, a loop closure could trigger the criterion after updating the map. Note that the map's accuracy is embedded in $\Delta\mathcal{U}$ given the graph-SLAM basis.

Should Γ drop below a certain threshold (Γ_{th}) during a given action window, w , active SLAM would be terminated as no information would be added to the system. Thus, for a fixed map size, if the robot keeps gathering the same information over and over (i.e. $\Delta\mathcal{U} \lesssim 0$), it will stop. Otherwise, exploration will continue even when information decreases (pure exploration) until the information loss exceeds the gain in the map area. This would indicate that a different exploration strategy is required, rather than that exploration has finished. Therefore, the proposed SC depends on two parameters: the minimum increment Γ_{th} , and the window in which evaluations are performed. The former is related to task-completion, and the smaller it is, the most complete, accurate and time-expensive the task will be. The latter serves to check whether it is a stationary

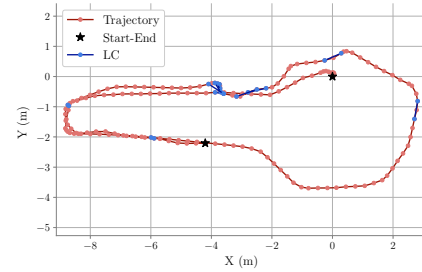


Fig. 1. Trajectory of the active SLAM example experiment.

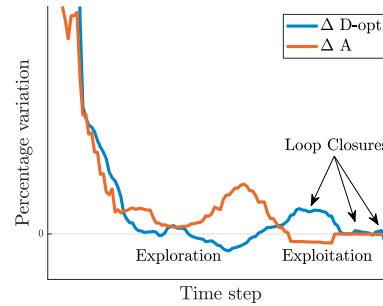


Fig. 2. Evolution of ΔD -opt and ΔA during the active SLAM example experiment.

or transient regime: short horizons (e.g. 1 action) should not be used to check for termination.

Let us consider a simple example of active SLAM to illustrate the typical trend of both terms. Figure 1 contains its trajectory (red dots correspond to graph nodes and red lines to odometry edges) and loop closures (blue). Figure 2 shows the percentage variation in D -optimality (blue) and area (red) during the task. Note that D -opt is computed over \mathbf{Y} and thus denotes knowledge over the robot's localisation. Two stages can be clearly distinguished, also labelled in the figure: exploration and exploitation. During the first stage, there is a large initial increase in the known area and localisation information. As exploration continues, new areas decrease and the robot's uncertainty gradually increases. At a certain point, the robot starts to revisit known places, entering the second stage. During exploitation, the area no longer increases, and may even decrease after closing a loop due to map updates. On the other hand, localisation uncertainty continues to decrease as loops are closed. However, once the first few loop closures occur, there is no further improvement. Both increments reach a constant regime near zero and after this, if the robot keeps gathering the same data from the environment, there is a loss of information (see final part of Figure 2). This regime represents the point at which exploration should be terminated or the utility function changed.

We are aware that evaluating variations on D -opt rather than raw values does not allow to directly constrain the robot's variance. Still, contrary to usual SC in the literature, it does neither imply prior knowledge of the environment nor human interaction during the task, takes into account both variables of interest and is transferable across systems.

Table 2. Results of exploration in AWS scenarios with agents with different SC. They have been ordered by the time consumed. † denotes values are explicitly fixed by the criterion.

Scenario	SC	Time (s)	Area (m ²)	Coverage (%)	mRMSE (m)	n	d	Opt.	D-opt
Bookstore	Temporal (10 min)	†600	155.04	86.01	0.22	492	2.48	3	3372.58
	<i>Ours</i>	1018	159.56	88.52	0.11	867	2.74	6	3862.94
	Coverage (90%)	∞	–	†90	–	–	–	–	–
	Coverage (99%)	∞	–	†99	–	–	–	–	–
House	Coverage (90%)	301	149.53	†90	0.72	236	2.49	0	2869.24
	<i>Ours</i>	482	155.40	93.53	0.23	331	2.45	1	2941.25
	Temporal (10 min)	†600	155.48	93.58	0.24	382	2.40	2	2909.94
	Coverage (99%)	∞	–	†99	–	–	–	–	–

4.1 Experimental Validation

In order to test the proposed criterion, active SLAM was performed in two scenarios in Gazebo, which are slightly modified versions of AWS bookstore and house worlds (e.g., they are now closed environments). See figures 3a and 3d. The robot is a wheeled platform equipped with a laser sensor with 180° field of view, 5 m range and 1500 beams in each scan. The maximum linear and angular velocities of the robot are 0.2 m s⁻¹ and 0.8 rad s⁻¹, respectively. Active SLAM framework is the same as in (Placed and Castellanos, 2021b), which is publicly available.

The proposed experiment consists of determining and comparing the moment at which different SC decide to stop exploration. To make a fair comparison, different criteria are evaluated during the same exploration run. That is, as the autonomous agent explores the environment, SC are checked and if any of them is met, results for that criterion are extracted. Then, the exploration is resumed. Apart from the task-driven criterion ($\Gamma_{th} = 2\%$, $w = 3$), we used a time-based one (10 min) and two coverage conditions: 90 and 99%, following the somewhat standardised values (Xu et al., 2022). Table 2 shows results (mean over 2 trials) of all agents in both scenarios. It contains the time elapsed until the stopping conditions were met, map metrics (area explored, coverage, max. RMSE) and graph metrics (size, average node degree, number of graph optimisations and final graph *D-opt*). Note that coverage conditions assume the environment known, and that the time has been arbitrarily selected. Fig. 3 shows the maps and pose-graphs at the moment of fulfilment of the different criteria.

In the bookstore environment, time-based criterion is the first to be met: after 10 minutes of exploration, the agent knows 86% of the scenario and maximum error of the mapped area is 22 cm. Almost 7 minutes later, the proposed criterion is satisfied. Despite the similar coverage (only 2.5% higher), the number of graph optimisations has doubled and map error has halved. Also, the robot localisation information has increased by 15%. Coverage conditions are never satisfied, as there are many unobservable or unreachable areas for the robot.

On the other hand, the 90% coverage criterion is the first to be fulfilled in the second scenario. However, at that time the accuracy of the robot localisation and of the built map are extremely low. In addition, many interesting areas have not yet been explored (see Figure 3). After 8 minutes of exploration, the TOED-based criterion is met: not only

3.5% more of the environment has been explored, but also known areas have been exploited: map maximum error has been reduced threefold and the localisation information has increased. Finally, after 10 minutes, the temporal criterion triggers, which, regarding the map and coverage, has very similar results to our criterion. This experiment illustrates how repeatedly acquiring the same data (over-exploitation) can be indeed detrimental for the SLAM algorithm: despite a higher number of loop closures, the robot localisation and the map are less accurate and the graph is not as well connected (see last columns of Table 2). Again, the 99% coverage criterion is never met.

The proposed task-driven criterion achieves, in all cases, a relevant behaviour for active SLAM, outperforming other studied criteria. On the one hand, coverage criteria require prior knowledge of the environment and, even then, may fail. On the other hand, time-based criteria either require prior knowledge or are impractical: large times may lead to over-exploitation while short may fail to visit areas of interest. Our criterion succeeds in autonomously identifying the moment in which exploration is no longer adding relevant information, without requiring prior knowledge or manual tuning beyond the specification of the two parameters it depends on.

5. CONCLUSIONS

In this manuscript, we have studied the often overlooked problem of autonomously deciding when robotic exploration should terminate. First, we have thoroughly reviewed the use of existing *stopping criteria* in the literature, highlighting the assumptions on which they are based and their limitations for performing active SLAM in real environments. Motivated by this, and following the recent work on fast computation of *optimality criteria* in active graph-SLAM, we have proposed a TOED-based *stopping criterion*, which autonomously decides when exploration is not adding further information to the SLAM system. Experimental results show the relevance of this criterion for active SLAM and how it outperforms others used in the literature. This seminal work aims to stimulate the use of task-driven *stopping criteria* in active SLAM, rather than coverage or temporal metrics.

Future work will aim to study the use of raw *optimality criteria* rather than their evolution, and will evaluate their use to decide to change the exploration strategy (i.e. switch the utility function) instead of just terminating exploration.

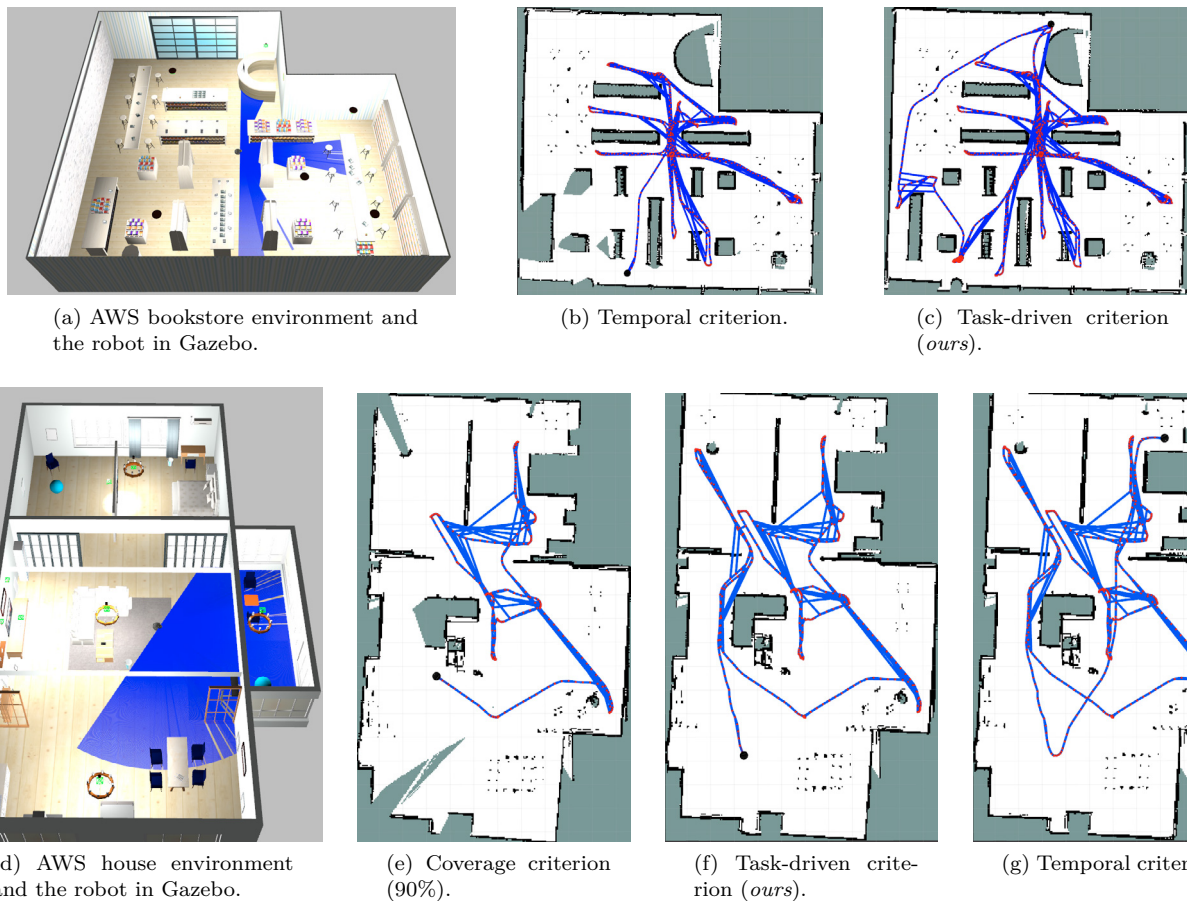


Fig. 3. View of the AWS bookstore (a) and house (d) environments in Gazebo. Also, the maps and pose-graphs built in one trial using temporal, coverage and TOED-based *stopping criteria* in both environments. Results for coverage SC have been omitted when unfeasible.

REFERENCES

- Amigoni, F., Li, A.Q., and Holz, D. (2013). Evaluating the impact of perception and decision timing on autonomous robotic exploration. In *2013 European Conference on Mobile Robots*, 68–73. IEEE.
- Bajcsy, R. (1988). Active perception. *Proceedings of the IEEE*, 76(8), 966–1005.
- Barraquand, J. and Ferbach, P. (1995). Motion planning with uncertainty: The information space approach. In *Proceedings of 1995 IEEE International Conference on Robotics and Automation*, volume 2, 1341–1348. IEEE.
- Bircher, A., Kamel, M., Alexis, K., Oleynikova, H., and Siegwart, R. (2018). Receding horizon path planning for 3D exploration and surface inspection. *Autonomous Robots*, 42(2), 291–306.
- Cadena, C., Carlone, L., Carrillo, H., Latif, Y., Scaramuzza, D., Neira, J., Reid, I., and Leonard, J.J. (2016). Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age. *IEEE Transactions on Robotics*, 32(6), 1309–1332.
- Carrillo, H., Dames, P., Kumar, V., and Castellanos, J.A. (2018). Autonomous robotic exploration using a utility function based on Rényi’s general theory of entropy. *Autonomous Robots*, 42(2), 235–256.
- Chen, F., Bai, S., Shan, T., and Englot, B. (2019). Self-learning exploration and mapping for mobile robots via deep reinforcement learning. In *AIAA SciTech Forum*.
- Chen, Y., Huang, S., and Fitch, R. (2020). Active SLAM for mobile robots with area coverage and obstacle avoidance. *IEEE/ASME Transactions on Mechatronics*.
- Deng, D., Duan, R., Liu, J., Sheng, K., and Shimada, K. (2020). Robotic exploration of unknown 2D environment using a frontier-based automatic-differentiable information gain measure. In *2020 IEEE/ASME International Conference on Advanced Intelligent Mechatronics (AIM)*, 1497–1503. IEEE.
- Durrant-Whyte, H. and Bailey, T. (2006). Simultaneous localization and mapping: part i. *IEEE Robotics & Automation Magazine*, 13(2), 99–110.
- Elimelech, K. and Indelman, V. (2019). Simplified decision making in the belief space using belief sparsification. *arXiv preprint arXiv:1909.00885*.
- Ghaffari Jadidi, M., Valls Miro, J., and Dissanayake, G. (2019). Sampling-based incremental information gathering with applications to robotic exploration and environmental monitoring. *The International Journal of Robotics Research*, 38(6), 658–685.
- Gomez, C., Hernandez, A.C., and Barber, R. (2019). Topological frontier-based exploration and map-building using semantic information. *Sensors*, 19(20), 4595.
- Grisetti, G., Kummerle, R., Stachniss, C., and Burgard, W. (2010). A tutorial on graph-based SLAM. *IEEE Intelligent Transportation Systems Magazine*, 2(4), 31–43.

- Indelman, V., Carlone, L., and Dellaert, F. (2015). Planning in the continuous domain: A generalized belief space approach for autonomous navigation in unknown environments. *The International Journal of Robotics Research*, 34(7), 849–882.
- Khosoussi, K., Giamou, M., Sukhatme, G.S., Huang, S., Dissanayake, G., and How, J.P. (2019). Reliable graphs for SLAM. *The International Journal of Robotics Research*, 38(2-3), 260–298.
- Khosoussi, K., Huang, S., and Dissanayake, G. (2014). Novel insights into the impact of graph structure on SLAM. In *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2707–2714. IEEE.
- Kitanov, A. and Indelman, V. (2019). Topological information-theoretic belief space planning with optimality guarantees. *arXiv preprint arXiv:1903.00927*.
- Korb, R. and Schöttl, A. (2018). Exploring unstructured environment with frontier trees. *Journal of Intelligent & Robotic Systems*, 91(3), 617–628.
- Lenac, K., Kitanov, A., Maurović, I., Dakulović, M., and Petrović, I. (2016). Fast active SLAM for accurate and complete coverage mapping of unknown environments. In *Intelligent Autonomous Systems 13*, 415–428. Springer.
- Leung, C., Huang, S., and Dissanayake, G. (2008). Active SLAM in structured environments. In *2008 IEEE International conference on Robotics and Automation*, 1898–1903. IEEE.
- Lluvia, I., Lazkano, E., and Ansuategi, A. (2021). Active mapping and robot exploration: A survey. *Sensors*, 21(7), 2445.
- Makarenko, A.A., Williams, S.B., Bourgault, F., and Durrant-Whyte, H.F. (2002). An experiment in integrated exploration. In *2002 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, volume 1, 534–539. IEEE.
- Montemerlo, M. and Thrun, S. (2003). Simultaneous localization and mapping with unknown data association using FastSLAM. In *2003 IEEE International Conference on Robotics and Automation (Cat. No. 03CH37422)*, volume 2, 1985–1991. IEEE.
- Palomeras, N., Carreras, M., and Andrade-Cetto, J. (2019). Active SLAM for autonomous underwater exploration. *Remote Sensing*, 11(23), 2827.
- Pham, V.C. and Juang, J.C. (2013). A multi-robot, cooperative, and active SLAM algorithm for exploration. *International Journal of Innovative Computing, Information and Control*, 9(6), 2567–2583.
- Pimentel, J.M., Alvim, M.S., Campos, M.F., and Macharet, D.G. (2018). Information-driven rapidly-exploring random tree for efficient environment exploration. *Journal of Intelligent & Robotic Systems*, 91(2), 313–331.
- Placed, J.A. and Castellanos, J.A. (2020). A deep reinforcement learning approach for active SLAM. *Applied Sciences*, 10(23), 8386.
- Placed, J.A. and Castellanos, J.A. (2021a). Fast autonomous robotic exploration using the underlying graph structure. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 6672–6679. IEEE.
- Placed, J.A. and Castellanos, J.A. (2021b). Fast uncertainty quantification for active graph SLAM. *arXiv preprint arXiv:2110.01289*.
- Platt, R., Tedrake, R., Kaelbling, L., and Lozano-Perez, T. (2010). Belief space planning assuming maximum likelihood observations. In *Proceedings of the Robotics: Science and Systems Conference, 6th, 2010*.
- Prentice, S. and Roy, N. (2009). The belief roadmap: Efficient planning in belief space by factoring the covariance. *The International Journal of Robotics Research*, 28(11-12), 1448–1465.
- Pukelsheim, F. (2006). *Optimal design of experiments*.
- Rodríguez-Arévalo, M.L., Neira, J., and Castellanos, J.A. (2018). On the importance of uncertainty representation in active SLAM. *IEEE Transactions on Robotics*, 34(3), 829–834.
- Salan, S., Drumwright, E., and Lin, K.I. (2014). Minimum-energy robotic exploration: A formulation and an approach. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 45(1), 175–182.
- Simmons, R., Apfelbaum, D., Burgard, W., Fox, D., Moors, M., Thrun, S., and Younes, H. (2000). Coordination for multi-robot exploration and mapping. In *Seventeenth AAAI Conference on Artificial Intelligence*, 852–858.
- Stachniss, C. (2009). *Robotic mapping and exploration*, volume 55. Springer.
- Stachniss, C. and Burgard, W. (2003). Mapping and exploration with mobile robots using coverage maps. In *2003 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, volume 1, 467–472. IEEE.
- Stachniss, C., Hahnel, D., and Burgard, W. (2004). Exploration with active loop-closing for FastSLAM. In *2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, volume 2, 1505–1510. IEEE.
- Suresh, S., Sodhi, P., Mangelson, J.G., Wettergreen, D., and Kaess, M. (2020). Active SLAM using 3D submap saliency for underwater volumetric exploration. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, 3132–3138. IEEE.
- Thrun, S., Burgard, W., and Fox, D. (2005). *Probabilistic Robotics*. MIT Press.
- Valencia, R., Valls Miro, J., Dissanayake, G., and Andrade-Cetto, J. (2012). Active pose SLAM. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*.
- Valencia, R. and Andrade-Cetto, J. (2018). Active pose SLAM. In *Mapping, Planning and Exploration with Pose SLAM*, 89–108. Springer.
- Whaite, P. and Ferrie, F.P. (1997). Autonomous exploration: Driven by uncertainty. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(3), 193–205.
- Xu, Y., Yu, J., Tang, J., Qiu, J., Wang, J., Shen, Y., Wang, Y., and Yang, H. (2022). Explore-bench: Data sets, metrics and evaluations for frontier-based and deep-reinforcement-learning-based autonomous exploration. *arXiv preprint arXiv:2202.11931*.
- Yamauchi, B., Schultz, A., and Adams, W. (1999). Integrating exploration and localization for mobile robots. *Adaptive Behavior*, 7(2), 217–229.