Some Performance Bounds of Strategies for Graph Exploration

(Extended Abstract)

Alessandro Riva
Politecnico di Milano
Piazza Leonardo da Vinci, 32
Milano, Italy
alessandro.riva10@mail.polimi.it

Alberto Quattrini Li
Politecnico di Milano
Piazza Leonardo da Vinci, 32
Milano, Italy
alberto.quattrini@polimi.it

Francesco Amigoni
Politecnico di Milano
Piazza Leonardo da Vinci, 32
Milano, Italy
francesco.amigoni@polimi.it

ABSTRACT

Exploration of unknown environments is relevant for many robotics applications, like map building and coverage. Several works in the literature have proposed exploration strategies that drive a mobile robot to greedily choose where to go next in order to incrementally map an initially unknown environment. In this paper, we theoretically study the worst and average traveled distance required to explore graph-based environments by some exploration strategies that consider distance and information gain in selecting the next destination location.

Categories and Subject Descriptors

I.2 [**Artificial Intelligence**]: Robotics

General Terms

Algorithms, Theory

Keywords

Graph exploration, robot exploration, online algorithms

1. INTRODUCTION

The mainstream approach to robot exploration of initially unknown environments is greedy [6], where candidate destination locations are usually selected on the *frontiers* between the known and the unknown portions of the environment [7], according to an *exploration strategy* that considers different criteria, like distance from the current position of the robot [7] and expected information gain of the candidate locations (e.g., [5]), in a utility function.

The assessment of such exploration strategies performed in the field of robotics is mainly empirical (e.g., [1, 2]). The computational geometry and the theoretical computer science communities have studied the exploration problem, but the derived bounds are often relative to specific, and sometimes not fully realistic, contexts (e.g., closed tours for graph exploration [3]).

To the best of our knowledge, very few works have considered more practical settings for deriving bounds on the quality of solutions produced by exploration strategies, prominently the work in [4] and [6]. In that approach, a single robot should explore all the vertices of an undirected graph, whose edges have unitary cost, with a sensor that allows the perception of the current vertex and of an arbitrary number of other vertices. Worst-case lower and upper bounds independent of the sensor range are given for the number of edges that a robot has to traverse to explore the whole graph.

In this paper, we first refine and complement the analysis of [4] and [6] by deriving some worst-case bounds on the number of edge traversals for some exploration strategies that use distance, information gain, and a combination of them as criteria. Our bounds explicitly embed the sensor range r and a termination criterion that prescribes the perception of a fraction $p \in (0, 1]$ of the vertices, as common in some robotics applications, such as search and rescue. Second, we address an average-case analysis on the performance of some exploration strategies in a class of graphs that model indoor environments, which, to the best of our knowledge, has never appeared in the literature.

2. EXPLORATION PROCESS MODEL

The environment is represented by an undirected, connected, unweighted, and finite graph $G = (V, E)$, where the vertices V correspond to the locations where an autonomous mobile robot can move and the edges E represent the direct connections between these locations (as in [6]).

The robot operates according to the following steps: (0) starting from an initial vertex v_0 having no *a priori* knowledge about the graph G; at a generic time step i, while being in v_i, (1) it perceives the surrounding environment generating P_i, that is the set of vertices within a finite range $r \in \mathbb{R}_{>0}$ from v_i; (2) it integrates the perceived data within the current knowledge about the environment obtaining the set of vertices V_i; (3) it reaches a vertex in F_i (which are the vertices still not perceived and neighbors of at least one vertex in V_i), chosen according to an exploration strategy $\mathscr{S}$, and starts again from (1). This process continues until a percentage $p \in (0, 1]$ of the vertices of G are perceived by the robot, namely until $\frac{|V_i|}{|V|} \geq p$. So, in the end, the robot follows a sequence of vertices $\mathscr{P} = \langle v_0, v_1, \cdots, v_k \rangle$, called the *exploration path*, composed of selected frontier vertices

Appears in: *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2015), Bordini, Elkind, Weiss, Yolum (eds.), May 4–8, 2015, Istanbul, Turkey.*

$v_{i+1} \in F_i$, with $0 \leq i < k$. Our perception model allows the robot to acquire knowledge about the incident edges of vertices $v \in P_i$ and to recognize whether there is an edge between two known vertices $v', v'' \in V_i$. We assume that the perceptions and the movements of the robot are error-free (i.e., deterministic). As a consequence, the robot perfectly knows its position in the environment.

We consider exploration strategies $\mathscr{S}$ that evaluate a candidate vertex $v \in F_i$ from the current position v_i adopting the following criteria:

- $d_i(v_i, v)$ is the geodesic distance between v_i and v in the graph induced by V_i on G, augmented with v and with the edges (in E) between v and vertices in V_i,
- $g(v, V_i)$ is the expected information gain at v, and is equal to the number of vertices the robot could perceive in v minus those already known.

We consider three exploration strategies:

- $\mathscr{S}_d$, which selects locations by simply minimizing the distance $d()$ (as for example in [6]),
- $\mathscr{S}_g$, which chooses candidate locations maximizing the information gain $g()$ (as, for example, in [1]),
- $\mathscr{S}_{dg}$, based on $\mathscr{S}_d$ but breaking ties favoring vertices with larger information gain $g()$.

In all the three cases, further ties are broken randomly with uniform probability.

3. WORST- AND AVERAGE-CASE RESULTS

Table 1 summarizes the worst-case bounds we derived for the three exploration strategies (due to space constraints, proofs are not reported).

$UB_{\mathscr{S}_d} = 2\lvert V\rvert \left(\ln \frac{2\lvert V\rvert+\lfloor r\rfloor(\lfloor r\rfloor-2)-7}{(\lfloor r\rfloor+1)^2} - \frac{\lvert V\rvert+\lfloor r\rfloor(\lfloor r\rfloor-2)-5}{(\lvert V\rvert-2)(\lfloor r\rfloor+1)} + 2\right)$	$LB_{\mathscr{S}_d} = \Omega\left(\frac{\log\lvert V\rvert-\log(\lfloor r\rfloor+1)^2}{\log\log\lvert V\rvert}\lvert V\rvert\right)$
$UB_{\mathscr{S}_g} = \left(\lvert V\rvert - \frac{\lvert V\rvert-1}{\lfloor r\rfloor+1}\right)\left(2\frac{\lvert V\rvert-1}{\lfloor r\rfloor+1} - 1\right)$	$LB_{\mathscr{S}_g} = \frac{\lfloor r\rfloor+1}{2}\left(\frac{\lvert V\rvert-\lfloor r\rfloor}{\lfloor r\rfloor+1}\right)\left(\frac{\lvert V\rvert-\lfloor r\rfloor}{\lfloor r\rfloor+1} - 1\right)$
$UB_{\mathscr{S}_{dg}} = UB_{\mathscr{S}_d}$	$LB_{\mathscr{S}_{dg}} = LB_{\mathscr{S}_d}$

Table 1: Worst-case upper (left) and lower (right) bounds on the number of edge traversals for $\mathscr{S}_d$, $\mathscr{S}_g$, and $\mathscr{S}_{dg}$, on any undirected, connected, unweighted, and finite graph $G = (V, E)$, given $p = 1$ and $r \in \mathbb{R}_{\geq 1}$.

$\mathscr{S}_g$ is the exploration strategy with the highest worst-case upper bound, while $\mathscr{S}_d$ and $\mathscr{S}_{dg}$ have the same worst-case upper bounds, namely, there is no gain in the worst case using, besides distance, information gain as evaluation criterion. This is in line, for example, with some results obtained in real (or realistically simulated) environments that suggest that sometimes using information gain in the exploration strategies does not shorten the paths for completely exploring the environments [2, 5].

Since our bounds, differently of those of [4] and [6]. explicitly consider the range r, it is possible to observe that the impact of increasing perception range r on the worst-case length of exploration is significant for small values of r, becoming less significant for large values of r. This holds for all exploration strategies. Despite the fact that they have the same asymptotic complexity of those of [4] and [6] that do not explicitly embed p and r, the bounds we found have lower actual values. More generally, our analysis shows that with increasing r, the exploration process is shortened,

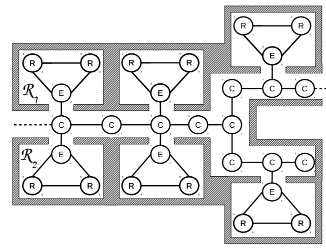

Figure 1: Example of graph that belongs to $\mathscr{G}$, vertices are labeled as C=corridor, R=room, and E=entrance.

which is an intuitively evident result consistent with several experimental findings (e.g., [1]).

Another insight is that, in the worst case, considering a percentage p of vertices to perceive has the same effect of scaling (by p) the number of vertices of the graph representing the environment. Given the worst-case bounds, an exploration strategy that considers only distance as criterion scales with p better than one that just considers information gain.

To distinguish the performance between $\mathscr{S}_d$ or $\mathscr{S}_{dg}$, we performed an average-case analysis and we obtained an estimate on the difference between the numbers of edges traversed by the robot using $\mathscr{S}_d$ or $\mathscr{S}_{dg}$ in a specific class $\mathscr{G}$ of graphs that model indoor environments (such as the one depicted in Figure 1). The result (not reported due to space limit) shows that, differently from the worst case, considering expected information gain in exploration strategies provides an advantage for graphs in $\mathscr{G}$. This can be intuitively explained as the robot visits all rooms encountered without the need to go back to visit some rooms left behind while traversing the corridor (see, e.g., Figure 1). This result is also supported by simulated experiments that we conducted in randomly generated environments that belongs to $\mathscr{G}$.

Our long-term goal is to aid the design of efficient exploration strategies from the insights provided by our analysis.

REFERENCES

[1] F. Amigoni. Experimental evaluation of some exploration strategies for mobile robots. In *Proc. ICRA*, pages 2818–2823, 2008.

[2] M. Julia, A. Gil, and Ó. Reinoso. A comparison of path planning strategies for autonomous exploration and mapping of unknown environments. *Auton Robot*, 33(4):427–444, 2012.

[3] B. Kalyanasundaram and K. Pruhs. Constructing competitive tours from local information. *Theor Comput Sci*, 130:125–138, 1994.

[4] S. Koenig, C. Tovey, and Y. Smirnov. Performance bounds for planning in unknown terrain. *Artif Intell*, 147(1-2):253–279, 2003.

[5] C. Stachniss and W. Burgard. Exploring unknown environments with mobile robots using coverage maps. In *Proc. IJCAI*, pages 1127–1134, 2003.

[6] C. Tovey and S. Koenig. Improved analysis of greedy mapping. In *Proc. IROS*, pages 3251–3257, 2003.

[7] B. Yamauchi. A frontier-based approach for autonomous exploration. In *Proc. CIRA*, pages 146–151, 1997.