Computers & Operations Research ∎ (∎∎∎∎) ∎∎∎–∎∎∎

Contents lists available at ScienceDirect

Computers & Operations Research

journal homepage: www.elsevier.com/locate/caor

A meta-heuristic based goal-selection strategy for mobile robot search in an unknown environment

Miroslav Kulich [a], Juan José Miranda-Bront [b,c,*], Libor Přeučil [a]

[a] Czech Institute of Informatics, Robotics, and Cybernetics, Czech Technical University in Prague, Czech Republic
[b] Departamento de Computación, FCEyN, Universidad de Buenos Aires, Pabellón I, Ciudad Universitaria, C1428EGA, CABA, Argentina
[c] Consejo Nacional de Investigaciones Científicas y Técnicas, Argentina

ARTICLE INFO

Article history:
Received 30 September 2015
Received in revised form
25 March 2016
Accepted 29 April 2016

Keywords:
Mobile robotics
Robotic search
Graph Search Problem
Traveling Deliveryman Problem
GRASP

ABSTRACT

The single-robot search problem in an unknown environment is defined as the problem of finding a stationary object in the environment whose map is not known a priori. Compared to exploration, the only difference lies in goal selection as the objectives of search and exploration are dissimilar, i.e. a trajectory that is optimal in exploration does not necessarily minimize the expected value of the time to find an object along it. For this reason, in this paper we extend the preliminary ideas presented in Kulich et al. [1] to a general framework that accounts for the particular characteristics of the search problem. Within this framework, an important decision involved in the determination of the trajectory can be formulated as an instance of the Graph Search Problem (GSP), a generalization of the well-known Traveling Deliveryman Problem (TDP) which has not received much attention in the literature. We developed a tailored Greedy Randomized Adaptive Search Procedure (GRASP) meta-heuristic for the GSP, which generates good quality solutions in very short computing times and is incorporated in the overall framework. The proposed approach produces very good results in a simulation environment, showing that it is feasible from a computational standpoint and the proposed strategy outperforms the standard approaches.

1. Introduction and literature review

Single-robot search, similar to exploration, can be understood as a process of autonomous navigation of a mobile robot in an unknown environment in order to find an object of interest. The search algorithm can be formulated as the iterative procedure consisting of model updating with actual sensory data, selection of a new goal for a robot based on the current knowledge of the environment, and subsequent navigation to this goal. A natural condition is to perform the search with an expected minimal usage of resources, e.g., trajectory length, time of search, or energy consumption. Robot search is an important task e.g. in the search and rescue scenario, where the goal is to find a black box flight recorder or debris after a plane crash, or victims/survivors after an accident or a catastrophe. In these situations, typically, the searched object does not move and a precise map of the environment is not available in advance.

While the research of exploration by a single or multiple mobile robots has been quite intensive (see e.g. [2–4], or our previous research [5,6]), the search problem has been addressed marginally by the robotic community. On the other hand, the general structure of the search problem is the same as that of the exploration problem. The only difference lies in goal selection as the objectives of search and exploration are dissimilar, i.e. a trajectory that is optimal in exploration does not necessarily minimize the expected value of the time to find an object along it.

Some effort has been devoted to the single and multi-robot search problem for a priori known environments which can be straightforwardly used as a goal selection strategy in the iterative procedure of the search problem in an unknown space. Sarmiento et al. [7] formulate the problem so that the time required to find an object is a random variable induced by a choice of search path and a uniform probability density function for the object's location. They propose two-stage process to solve the problem. Firstly, a set of locations (known as guards from the Art Gallery Problem [8]) to be visited is determined. An order of visiting those locations minimizing the expected time to find an object is found then. The optimal order is determined by a greedy algorithm in a reduced search space, which computes a utility function for several steps

* Corresponding author at: Departamento de Computación, FCEyN, Universidad de Buenos Aires, Pabellón I, Ciudad Universitaria, C1428EGA, CABA, Argentina.
E-mail addresses: kulich@ciirc.cvut.cz (M. Kulich), jmiranda@dc.uba.ar (J.J. Miranda-Bront), preucil@ciirc.cvut.cz (L. Přeučil).

http://dx.doi.org/10.1016/j.cor.2016.04.029
0305-0548/

ahead. This approach is then used in Sarmiento et al. [9], where robot control is assumed with the aim to generate smooth and locally optimal trajectories. Hollinger et al. [10] utilize a Bayesian network for estimating the posterior distribution of target's position and present a graph search to minimize the expected time needed to capture a non-adversarial object.

In the operations research literature, the single-robot search problem in known environments is formulated as the Graph Search Problem (GSP), introduced in Kutsoupias et al. [11] and arising in applications to web searching problems. The GSP can be defined as follows: consider a complete graph $G = (V, E)$, with $V = \{0, 1, \dots, n\}$, a distance function d_{ij}, $(i, j) \in E$, and a probability p_i that the required information is at vertex $i \in V$. The objective of the GSP is to find a tour that minimizes the expected time to find the required information. Besides some theoretical results regarding approximation schemes presented in Ausiello et al. [12], no further developments are present in the related literature.

The Traveling Deliveryman Problem (TDP) is a well known problem in the operational research community, which has received some attention in the last few years. It can be formulated under the same settings as the GSP, with the only difference that the probability of finding the information is the same for all the vertices in the graph. Integer Linear Programming (ILP) formulations and exact algorithms for the TDP are considered in, e.g., [13–15]. Recently, several ILP formulations and Branch-and-Cut (BC) and Branch-Cut-and-Price (BCP) algorithms have been developed for the Time-Dependent Traveling Salesman Problem (TDTSP), a generalization of the classical TSP which also generalizes the TDP. Some of these approaches are proposed in [16–19]. Overall, the best exact algorithm is the BCP proposed in Abeledo et al. [16], being able to solve instances with up to 107 vertices to optimality in several hours of computing time.

In addition, several heuristics and meta-heuristics have been proposed for the TDP and some other variants. The approaches rely mostly on Greedy Randomized Adaptive Search Procedure (GRASP), introduced originally by Feo and Resende [20], and Variable Neighborhood Search (VNS), proposed by Hansen and Mladenovic [21]. Salehipour et al. [22] propose a GRASP, evaluating the impact of considering for the local search phase a Variable Neighborhood Descent (VND) as well as a VNS procedure. Mladenovic et al. [23] propose a General VNS (GVNS), which is able to improve the results obtained by Salehipour et al. [22]. Silva et al. [24] propose a simple multistart heuristic combined with an Iterated Local Search procedure. The method improves all the results reported in Salehipour et al. [22] and finds a new best known solution in two benchmark instances. To the best of our knowledge, the approach by Silva et al. [24] is the one producing the best results in the literature. Finally, regarding variants of the TDP, Dewilde et al. [25] tackle the TDP with profits and Heilporn et al. [26] the TDP with time windows.

The single-robot search problem in an unknown environment is formally formulated in our previous research Kulich et al. [1], where a criterion to be optimized is defined and several goal-selection strategies are considered. Two of these strategies are borrowed from the exploration problem, for which they were originally designed, namely the *greedy strategy* proposed in Yamauchi [2] and the *Traveling Salesman strategy* introduced in Kulich et al. [5], while the third one is based on the formulation and resolution of a TDP instance within the framework. The three strategies are evaluated in a simulation environment and the behavior is discussed. The results are somehow mixed, showing that none of these strategies clearly dominates the others. As the key problem was identified that the studied strategies do not use information gain of visiting a goal, which may be crucial for designing effective search strategies.

The aim of this research goes in that direction. The contribution of this paper is threefold. Firstly, we build upon the preliminary results presented in Kulich [1] and provide a complete framework that accounts for the particular characteristics of the problem. Secondly, in order to provide a complete implementation of such framework, we develop a meta-heuristic to solve the GSP, tailored for the context of the search problem and employ it as a goal-selection strategy within the search framework. In the proposed framework, instances consider generally between 50 and 100 vertices, but the amount of time required to obtain a near-optimal solution should not exceed one second. This is a key factor since, opposed to the TSP, even the TDP requires several hours of execution to solve to optimality instances of moderate size. Finally, the whole framework is evaluated computationally in a simulation environment, where the proposed goal selection strategy outperforms the traditional one and shows that the proposed strategy has potential to be applied in practice.

The rest of the paper is organized as follows. The problem definition is presented in Section 2, while the frontier-based framework for search is described in Section 3. The GRASP approach developed for the GSP is described in Section 4. In Section 5 we present the computational results, including the evaluation of the GRASP meta-heuristic and the evaluation of the overall framework in a simulation environment. Finally, Section 6 is dedicated to concluding remarks and future directions.

2. Problem definition

Assume an autonomous mobile robot equipped with a ranging sensor with a fixed, limited range (e.g. laser range-finder) and 360° field of view operating in a priori unknown environment. The search problem is defined as the process of navigation of the robot through this environment with the aim to find a stationary object of interest placed in the environment randomly and reachable by the robot.[1] By finding an object we understand the situation when it is firstly detected by robot's sensors. A natural condition is to minimize the time when this situation occurs. More formally, this condition can be expressed as minimization of the expected (mean) time T_f the object is firstly detected, when the robot follows the trajectory R:

$$T_f = \mathbb{E}(T|R) = \int_0^\infty tp(t)\,\mathrm{d}t, \tag{1}$$

where $p(t)$ is the probability of finding the object at time t.

Sensing as well as planning is performed in discrete times, therefore (1) can be rewritten as

$$T_f = \mathbb{E}(T|R) = \sum_{t=0}^{\infty} tp(t), \tag{2}$$

where $p(t) = \frac{A_t^R}{A_{total}}$ is the ratio of the area A_t^R newly sensed at time t when the robot follows the trajectory R and A_{total} the area of the whole environment the robot operates. The objective is to find the trajectory R^{opt} minimizing (2):

$$R^{opt} = \arg\min_R \mathbb{E}(T|R) = \arg\min_R \sum_{t=0}^{\infty} tA_t^R. \tag{3}$$

Notice that A_{total} can be omitted as it is a constant.

[1] In general, a priori information about object's position can be given in the form of a probability density function (PDF), but a uniform distribution is expected in the paper. On the other hand, incorporation of a priori PDF is straightforward as it involves only determination of probabilities $p(t)$ in (2).