Artificial Intelligence 196 (2013) 26–52

Contents lists available at SciVerse ScienceDirect

Artificial Intelligence

www.elsevier.com/locate/artint

Physical search problems with probabilistic knowledge ☆

Noam Hazon [a,∗,1], Yonatan Aumann [b], Sarit Kraus [b], David Sarne [b]

[a] *Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, USA*
[b] *Department of Computer Science, Bar-Ilan University, Ramat Gan, Israel*

A R T I C L E I N F O

Article history:
Received 31 August 2011
Received in revised form 30 September 2012
Accepted 24 December 2012
Available online 3 January 2013

Keywords:
Graph search
Economic search

A B S T R A C T

This paper considers the problem of an agent or a team of agents searching for a resource or tangible good in a physical environment, where the resource or good may possibly be obtained at one of several locations. The cost of acquiring the resource or good at a given location is uncertain (a priori), and the agents can observe the true cost only when physically arriving at this location. Sample applications include agents in exploration and patrol missions (e.g., an agent seeking to find the best location to deploy sensing equipment along its path). The uniqueness of these settings is in that the cost of observing a new location is determined by distance from the current one, impacting the consideration for the optimal search order. Although this model captures many real world scenarios, it has not been investigated so far.

We analyze three variants of the problem, differing in their objective: minimizing the total expected cost, maximizing the success probability given an initial budget, and minimizing the budget necessary to obtain a given success probability. For each variant, we first introduce and analyze the problem with a single agent, either providing a polynomial solution to the problem or proving it is NP-complete. We also introduce a fully polynomial time approximation scheme algorithm for the minimum budget variant. In the multi-agent case, we analyze two models for managing resources, shared and private budget models. We present polynomial algorithms that work for any fixed number of agents, in the shared or private budget model. For non-communicating agents in the private budget model, we present a polynomial algorithm that is suitable for any number of agents. We also analyze the difference between homogeneous and heterogeneous agents, both with respect to their allotted resources and with respect to their capabilities. Finally, we define our problem in an environment with self-interested agents. We show how to find a Nash equilibrium in polynomial time, and prove that the bound on the performance of our algorithms, with respect to the social welfare, is tight.

1. Introduction

Frequently, in order to successfully complete its task, an agent may need to *explore* (i.e., search) its environment and choose among different available options. For example, an agent seeking to purchase a product over the Internet needs to query several electronic merchants in order to learn their posted prices; a robot searching for a resource or a tangible good needs to travel to possible locations where the resource is available and learn the configuration in which it is available, as

☆ This paper extends two earlier conference papers (Aumann et al., 2008 [6]; Hazon et al., 2009 [31]).

∗ Corresponding author.
E-mail addresses: noamh@cs.cmu.edu (N. Hazon), aumann@cs.biu.ac.il (Y. Aumann), sarit@cs.biu.ac.il (S. Kraus), sarned@cs.biu.ac.il (D. Sarne).

[1] This work was done while the author was at Bar-Ilan University.

0004-3702/$ – see front matter
http://dx.doi.org/10.1016/j.artint.2012.12.003

well as the difficulty of obtaining it there. In these environments, the benefit associated with an opportunity is revealed only upon observing it. The only knowledge available to the agent prior to observing the opportunity is the probability associated with each possible value of each prospect.

While in virtual environments the exploration can sometimes be considered costless, in physical environments traveling and observing typically entails a cost. Furthermore, traveling to a new location may increase or decrease the distance to other locations, so the cost associated with exploring other unexplored locations changes. For example, consider a Rover robot with the goal of mining a certain mineral. Potential mining locations may be identified based on satellite imaging, each location associated with some uncertainty regarding the difficulty of mining there. In order to assess the amount of battery power required for mining at a specific location, the robot needs to physically visit there. The robot's battery is thus used not only for mining the mineral but also for traveling from one potential location to another. Consequently, an agent's strategy in an environment associated with search costs should maximize the *overall* benefit resulting from the search process, defined as the value of the option eventually used, minus the costs accumulated along the process, rather than merely finding the best valued option.

In physical environments, it is common to use a team of agents rather than a single agent. Extending the single agent solution to multi-agent strategy may require subdividing the search space among the different agents. However, if agents have means of communication, then they may not wish to become too distant, as they can call upon each other for assistance. For example, even if a Rover does not have sufficient battery power for mining at a given location, it may be useful for it to travel to the site in order to determine the exact mining cost, and call for other robots that do have the necessary battery power. In this case, the scheduling of the robots' travel times is key, and must be carefully planned. If the agents are not fully cooperative, a selfish behavior should also be considered. Each one of the agents will try to minimize its traveling costs while still achieving the group's goal.

Finally, agents may be of different types, or with different amounts of resources. For example, Rover robots may be entering the mission with differing initial battery charges. They may also differ in their capabilities, like a team of Rovers in which some were specifically designed for mining missions, and thus require less battery power for the same mining task.

This paper aims at taking the first steps in understanding the characteristics of such physical search environments, both for the single and multi agent cases, and developing efficient exploration strategies for the like. Our main focus is on the case where the opportunities are aligned along a path, as in the case of perimeter patrol [60,19,2,3]. We note that many single and multi-agent coverage algorithms convert their complex environment into a simple long path [52,25,32]. Furthermore, we show that the problem in more general metric spaces is NP-complete, even for a tree graphs. For exposition purposes, in the remainder of the paper we use the classical procurement application where the goal of the search is purchasing a product and the value of each observed opportunity represents a price. Of course, this is only one example of the general setting of exploration in a physical environment, and the discussion and results of this paper are relevant to any such setting, provided that exploration and fulfilling the task consume the same type of resource.

We consider three variants of the problem, differing in their objective. The first (*Min-Expected-Cost*) is the problem of an agent that aims to minimize the expected total cost of completing its task. The second (*Max-Probability*) considers an agent that is given a budget for the task (which it cannot exceed) and aims to maximize the probability it will complete the task (e.g., reach at least one opportunity with a budget large enough to successfully buy the product). In the last variant (*Min-Budget*) the agent is required to guarantee a pre-defined probability of completing the task, and aims to minimize the overall budget that will be required to achieve the said success probability. We also consider the multi-agent extensions of these variants. While the first variant fits mostly product procurement applications, the two latter variants fit well into applications of robots engaged in remote exploration, operating with a limited amount of battery power (i.e., a budget).

1.1. Summary of results

We first consider the single agent case. We prove that in general metric spaces all three problem variants are NP-hard. Thus, as mentioned, we focus on the setting where all locations are located along a path. For this setting we provide polynomial algorithms for the *Min-Expected-Cost* problem. We show the other two problems (*Min-Budget* and *Max-Probability*) to be NP-complete even for the path. Thus, we consider further restrictions and also provide an approximation scheme. We show that both problems are polynomial if the number of possible prices is constant. Even with this restriction, we show that these problems are NP-complete on a tree graph. For the *Min-Budget* problem, we provide an FPTAS (fully-polynomial-time-approximation-scheme), that provides a $(1+\epsilon)$ approximation for any $\epsilon > 0$, in time $O(\text{poly}(n\epsilon^{-1}))$, where n is the size of the input.

For the multi-agent case, we first analyze a shared budget model, where all the resources and costs are shared among all the agents. We show that if the number of agents is fixed, then all of the single-agent algorithms extend to k-agents, with the time bounds growing exponentially in k. Therefore the computation of the agents' strategies can be performed whenever the number of agents is relatively moderate, a common scenario in many physical environments where several agents cooperate in exploration and search. If the number of agents is part of the input then the multi-agent versions of *Min-Budget* and *Max-Probability* are NP-complete even on the path and even with a single price.

We then investigate a model of private budgets, where each agent has its own initial budget. We again assume that the number of possible prices is bounded. In this case, we separately consider the setting where agents can communicate and the setting where they cannot. For non-communicating agents we show a polynomial algorithm for the *Max-Probability*