Probabilistic Physical Search on General Graphs: Approximations and Heuristics*

JAAMAS Track

Noam Hazon
Ariel University
Ariel, Israel
noamh@ariel.ac.il

Mira Gonen
Ariel University
Ariel, Israel
mirag@ariel.ac.il

KEYWORDS

graph search; planning under uncertainty

ACM Reference Format:
Noam Hazon and Mira Gonen. 2020. Probabilistic Physical Search on General Graphs: Approximations and Heuristics. In *Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020), Auckland, New Zealand, May 9–13, 2020,* IFAAMAS, 3 pages.

1 INTRODUCTION

An autonomous intelligent agent often needs to explore its environment and choose among different available alternatives. In many physical environments the exploration is costly, and the agent also faces uncertainty regarding the price of the possible alternatives. For example, consider a traveling purchaser seeking to obtain an item [7]. While there may be prior knowledge regarding candidate stores (e.g., based on search history), the actual price at any given site may only be determined upon reaching the site. In another domain, consider a Rover robot seeking to mine a certain mineral on the face of Mars [8, 9]. While there may be prior knowledge regarding candidate mining sites (e.g., based on satellite images) [3, 4], the actual cost associated with the mining at any given location, e.g., in terms of battery consumption, may depend on the exact conditions at each site (e.g., soil type, terrain, etc.), and hence are fully known only upon reaching the site.

These scenarios are referred to as probabilistic physical search problems, since there is a prior probabilistic knowledge regarding the price of the possible alternatives at each site, and traveling for the purpose of observing a price typically entails a cost. Furthermore, exploration and obtaining the item results in the expenditure of the same type of resource. The purchaser's money is used not only to obtain the item but also for traveling from one potential store to another; the robot's battery is used not only for mining the mineral but also for traveling from one potential location to another. Thus, the agent needs to carefully plan its exploration and balance its use of the available budget between the exploration cost and the purchasing cost.

We focus on the development of efficient exploration strategies for probabilistic physical search problems on graphs. The analysis of such problems was initiated by Aumann et al. [1], who showed that it is (computationally) hard to find the optimal solution on general graphs. Accordingly, they provide a thorough analysis of physical search problems on one-dimensional path graphs, both for single and multi-agent settings. However, many real-world physical environments may only be represented by two-dimensional graphs. For example, the Mars rover can freely move directly from any possible mining location to another (with an associated travel cost), while in path graphs the robot is restricted to move only to the two adjacent neighbors of its current location. Our work thus handles probabilistic physical search problems on general graphs. To the best of our knowledge, ours is the first to do this.

We consider two variants of the problem. The first variant, coined *Max-Probability*, considers an agent that is given an initial budget for the task (which it cannot exceed) and needs to act in a way that maximizes the probability it will complete its task (e.g., reach at least one opportunity with a budget large enough to successfully buy the product). In the second variant, coined *Min-Budget*, we are required to guarantee some pre-determined success probability, and the goal is to minimize the initial budget necessary in order to achieve the said success probability.

Of course, the procurement application and the Mars rover scenario are only two examples of the general setting of costly exploration in an uncertain physical environment, and the discussion and results of this paper are relevant to any such domain. For example, a search-and-rescue plane collecting debris from a missing airliner at various locations will want to optimize its chance of success given limited fuel while minimizing the risk of adverse weather. A UAV that has to loiter to classify an object with high confidence via imagery will want to choose the best spot that will guarantee the completion of the mission while minimizing the required battery. Since previous work showed that probabilistic physical search problems are hard on general graphs, we either need to consider approximations with guaranteed bounds or heuristics with practical running time. We do both.

2 THEORETICAL RESULTS

We first establish an interesting connection between *Max-Probability* and the *Deadline-TSP* problems [2], and as a result we are able to provide an $O(\log n)$ approximation for the former (where n is the number of sites), based on an $O(\log n)$ approximation for the latter, with the only requirement that the probabilities are not too small. This result also resolves two open problems of [1].

THEOREM 2.1. *The* Max-Probability *problem can be approximated within a ratio of* $O(\log n + \log k)$*, for any instance of the problem for which it holds that* $\frac{p_v(c_i)}{1-\sum_{j=1}^{i-1} p_v(c_j)} \geq 1/c$ *for every vertex* v *and any*

*This paper is an extended abstract. For the full paper, see [5]

Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020), B. An, N. Yorke-Smith, A. El Fallah Seghrouchni, G. Sukthankar (eds.), May 9–13, 2020, Auckland, New Zealand.

cost $c_i \in C_v$, where c is any constant larger than 1. For $k = O(n)$ the Max-Probability *problem can be approximated within a ratio of $O(\log n)$, for the same instances.*

Now consider the dual of the *Max-Probability* problem, the *Min-Budget* problem:

Theorem 2.2. *The* Min-Budget *problem with a specific instance of equal vertex costs and equal single probabilities can be approximated within a ratio of* $5 + \epsilon$, *for any* $\epsilon > 0$.

Theorem 2.3. *The Min-Budget problem is hard to approximate within a ratio of* $\alpha = 220/219 - \epsilon$, *for any* $1/219 > \epsilon > 0$.

We note that probabilistic physical search problems are not simply variants of the *Deadline-TSP* problem. Indeed, Theorem 2.2 shows that *Min-Budget* can sometimes be approximated, while the dual of the *Deadline-TSP* problem, namely the *Min-Length-Deadline-TSP*, is hard to approximate within any constant ratio:

Theorem 2.4. *The* Min-Length-Deadline-TSP *problem is hard to approximate within any constant ratio.*

3 EXPERIMENTAL RESUTLS

We now consider heuristics for practical use. We adapt a Greedy approach and an Ant Colony Optimization (ACO) algorithm to our setting. We also test two restricted variants of the optimal algorithm: the Bounded-Length (BL) heuristic, which bounds the solution's length, and the No-Backtrack (NB) heuristic, which only checks paths with bounded length and without repetitions. For the empirical evaluation of our heuristics we used a real graph structure with the traveling costs set as the real distance between the vertices, which was extracted from the GIS data of the highways network of the USA. In addition, we uniformly sampled 100 vertices from the graph as the possible initial locations for all of our algorithms. The probabilities for each vertex were generated using a Dirichlet distribution [6]. Some of our results are presented in Figures 1-4.

For the *Max-Probability* problem, BL found the success probabilities that were almost as good as the success probabilities found by the optimal algorithm (on average), and NB was only a little behind. Both heuristics returned solutions that were at most 90% from the optimal solutions. Greedy performed significantly worse, and ACO was only a little better than Greedy. However, when examining the corresponding running time NB and Greedy performed much better than BL and ACO.

For the *Min-Budget* problem, BL and NB required initial budgets that were at most 20% (on average) more than the optimal initial budgets. ACO required larger initial budgets, and Greedy performed worse. When examining the corresponding running time NB and Greedy performed again better than BL and ACO.

We conclude from our experimental analysis that NB clearly is the winner, for both *Max-Probability* and *Min-Budget* problems, since it runs very fast but still finds near-optimal solutions. We suggest considering a similar approach, i.e., performing a complete search that only checks paths with bounded length and without repetitions, for solving other planning problems on non-complete graphs in practice.

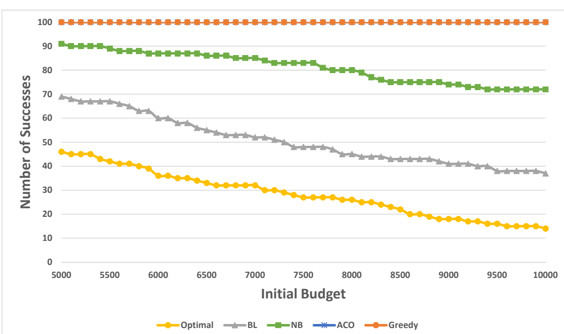

Figure 1: *Max-Probability*: the number of initial locations where a solution was found in less than 8 minutes.

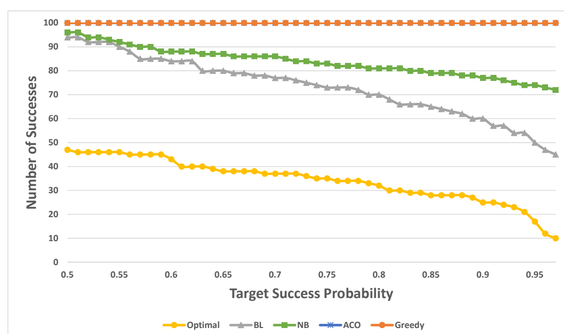

Figure 3: *Min-Budget*: the number of initial locations where a solution was found in less than 15 minutes.

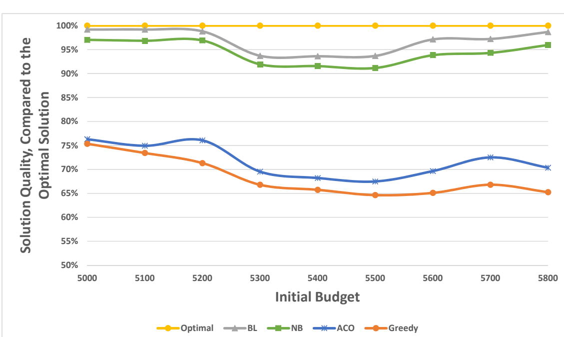

Figure 2: *Max-Probability*: performance comparison against the optimal algorithm. Higher is better.

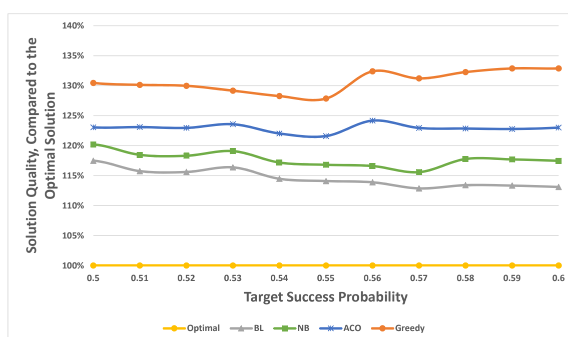

Figure 4: *Min-Budget*: performance comparison against the optimal algorithm. Lower is better.

REFERENCES

[1] Y. Aumann, N. Hazon, S. Kraus, and D. Sarne. 2008. Physical Search Problems Applying Economic Search Models. In *Proceedings of the Twenty-Third AAAI Conference on Artificial Intelligence (AAAI-2008)*. 9–16.

[2] N. Bansal, A. Blum, S. Chawla, and A. Meyerson. 2004. Approximation Algorithms for Deadline-TSP and Vehicle Routing with Time-Windows. In *Proceedings of the Thirty-Sixth Annual ACM Symposium on Theory of Computing (STOC-2004)*. 166–174.

[3] P. R. Christensen. 1986. The spatial distribution of rocks on mars. *ICARUS* 68 (1986), 217–238.

[4] M. P. Golombek, A. F. C. Haldemann, N. K. Forsberg-Taylor, E. N. DiMaggio, R. D. Schroeder, B. M. Jakosky, M. T. Mellon, and J. R. Matijevic. 2003. Rock size-frequency distributions on Mars and implications for Mars Exploration Rover landing safety and operations. *Journal of Geophysical Research: Planets* 108, E12 (2003).

[5] Noam Hazon and Mira Gonen. 2020. Probabilistic physical search on general graphs: approximations and heuristics. *Autonomous Agents and Multi-Agent Systems* 34, 1 (2020).

[6] Samuel Kotz, Narayanaswamy Balakrishnan, and Norman L Johnson. 2004. *Continuous multivariate distributions, Volume 1: Models and applications*. John Wiley & Sons, Chapter 49.

[7] T. Ramesh. 1981. Traveling purchaser problem. *Opsearch* 18 (1981), 78–91.

[8] D. Thompson, D. Wettergreen, G. Foil, M. Furlong, and A. Kiran. 2015. Spatio-Spectral Exploration Combining In Situ and Remote Measurements. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence (AAAI-2015)*. 3679–3685.

[9] D. Wettergreen, G. Foil, M. Furlong, and D. R. Thompson. 2014. Science autonomy for rover subsurface exploration of the atacama desert. *AI Magazine* 35, 4 (2014), 47–60.