The Role of Target Modeling in Designing Search Strategies

Alessandro Renzaglia, Narges Noori and Volkan Isler
University of Minnesota, Minneapolis, MN, USA

***Abstract*— This paper studies the problem of searching for an unknown moving target in a bounded two-dimensional convex area with a mobile robot. A key component of designing a search strategy is the target motion model, which is often unknown in practical scenarios. When designing search strategies, researchers either (1) ignore the target motion and treat the target as a stationary object with unknown location, (2) treat the target as an adversary and model the search task as a game, or (3) use a stochastic model such as a random walk. For each of these models we analyze possible search paths with the objective of minimizing the expected capture time. Our intent is to investigate how the choice of the model influences the choice of the strategy and consequently how the capture time will depend on this choice. In addition to a theoretical analysis, we compare the strategies in simulation.**

I. INTRODUCTION

Search problems are fundamental in robotics. Finding a target of unknown location in a given environment is a crucial task in many scenarios, such as finding victims in disaster areas, searching for intruders in restricted zones, tracking animals for conservation programs and so on. In many of these cases, not only the location of the target, but also its motion can be unknown to the searcher and a model has to be adopted. This makes the task of designing optimal strategies with the goal to reduce the time of finding the target even more challenging. A viable approach for this problem is to treat the target as an adversary which actively plans its path in order to avoid the searcher. Even though this scenario is not always realistic, it can provide a useful worst-case scenario for the original problem. On the other hand, this strategy can be very conservative and, depending on the capabilities of the target, an optimal strategy which assures capture may not exist. Furthermore, in some cases where the target is clearly non-adversarial, reasonable probabilistic models for the target's motion law can be considered. In this way, a more efficient strategy can be designed and the expected capture time can be significantly reduced. With this formulation, the problem is referred to as *one-sided probabilistic search* for a mobile target, where the target cannot observe the searcher and does not actively evade detection [7].

In our previous works [19], [16] we studied the problem of finding a target moving as a random walk in a one-dimensional environment. We approached the problem both with an analytical analysis for an energy-constrained searcher [19] and by formulating the problem as a Partially Observable Markov Decision Process (POMDP) considering a limited time window for the search [16]. In these works we showed how, without knowing the target's position and without sensing it until the capture, the probabilistic information can be exploited to design optimal paths which significantly outperform standard searching strategies.

In this paper, we extend this analysis to two-dimensional environments, which are more common in real-world applications. To simplify the analysis, we focus our attention on square environments, even though most of the presented results can be easily extended to any convex domain. We consider three possible ways of modeling the target's motion: stationary, adversarial and random walk motion.

Our contributions are as follows: We present a review of existing results in search problems and propose new strategies for various target models. For the adversarial case, the corresponding game is known as the Princess-and-Monster Game. This game was proposed by Isaacs [12] and solved by Gal [10]. We present an adaptation of Gal's strategy to square environments which is simple to analyze. In the case of a target moving as a random walk, we show how the probabilistic information can be used to design an adaptive search path. We also discuss the choice of possible searcher strategies in terms of expected energy required to accomplish the mission. Indeed, because of the diffusive properties of a random-walking target, the searcher can consider staying in a fixed location as a viable option. Finally, focusing on the case of a random walking target, we test and compare in simulation the performance of the proposed strategies in terms of mean capture time and capture probability. Even though the optimal search strategy for finding a two-dimensional random walker is still an open problem, simulations allow us to identify search strategies whose performance is close to known bounds. By simulating the search process, we also provide a comparison between stationary and moving search strategies, assuming that stay and move actions correspond to different energy costs.

The rest of the paper is organized as follows. We present an overview of related work in Section II. The search problems we aim to study are formalized in Section III. The capture strategies for various models are presented and discussed in Section IV, and the analysis of the capture probabilities in simulation is provided in Section V. Finally, we present concluding remarks and future work in Section VI.

II. RELATED WORK

Due to its great relevance in robotics applications, optimal search problems have been widely studied in past years. A

The authors are with the Department of Computer Science & Engineering, University of Minnesota, Minneapolis, USA.
Emails: {arenz, noori, isler}@cs.umn.edu
This work was supported by National Science Foundation Awards #1111638 and #0917676.

recent survey on results in autonomous search and pursuit-evasion problems, with a particular focus on applications in mobile robotic systems, can be found in [7]. In [8], the authors provide a probabilistic formulation for the search problem where a mobile searching agent seeks to locate a stationary target in a given search region or declare that the target is absent.

A possible POMDP-based solution for the problem of optimal searching for a non-adversarial target in an indoor environment by using a team of robots has been recently presented by Hollinger et al. in [11]. The authors assume to know the environment and choose searcher paths most likely to intersect with the path taken by the target. In this case the target's motion model is assumed to be known to the searcher.

Random motions, both as discrete random walks and continuous diffusive motions, have been extensively studied as models of animal motion or complex physical processes [13], [5]. In particular they are widely used in the literature to simplify pursuit-evasion games and absorption (or search) processes. A large number of interesting properties closely related to searching missions are collected in [17] including: first passage probability (the probability for the random walker of visiting for the first time a given point at a given time), survival probability (the probability that the random walk has not been found at a given time) and mean capture time (the expected time to be found). Various characteristics of random walks in general graphs have been studied in [15]. Examples are hitting time, which is the expected number of steps before a node is visited, and cover time, which is the expected number of steps to visit every node at least once. The survival probability of a particle that performs a random walk on a chain where traps are uniformly distributed with known concentration is studied in [3] and an asymptotically exact solution is provided. A similar problem, but in a grid with a single trap is studied in [21]. In [18], the authors study the survival probability of a prey on a line which is chased by more than one diffusive predator. The same problem but in a semi-infinite line where the boundary represents a haven for the prey is presented in [9]. Finally, in [14] the expected time required for capturing an adversarial robber and the random walk one is compared. Upper and lower bounds are derived for the ratio between these two values when the search is done on special graph structures.

The pursuit-evasion game corresponding to the adversarial version is known as the Princess and Monster Game. In this game, the players are in a dark room (i.e. they can not observe each other unless the monster captures the princess.) This game was proposed by Isaacs [12] and solved by Gal [10] who presented a randomized strategy to find the evader in time proportional to the area of the environment. In Section IV-B, we present an adaptation of this strategy for square environment which yields a simplified analysis. When the game takes place on a graph, it is known as the hunter and rabbit game. Adler et al. showed that the hunter can capture the rabbit in $O(N \log N)$ time in a graph with N vertices [1].

Fig. 1. Two ASVs used for the fish tracking project which motivates this work.

III. Problem Formulation

Let Ω be a planar bounded domain in which a target and a searcher can move. For simplicity we assume the environment to be a square of side L. Let the target's initial location $x_0 \in \Omega$ be chosen according to a probability distribution P_0. The searcher has a limited sensing range R_s, which we assume to be $R_s \ll L$. The detection of the target, or capture, will happen when

$$\|x_t - s_t\| < R_s\,, \tag{1}$$

where $\|\cdot\|$ represents the standard Euclidean norm and x_t, s_t are the target and searcher positions at time t respectively. The searcher can move within the search space with a velocity v. By choosing appropriate units, this quantity can be taken to be one. The objective is to design a searcher strategy which minimizes the expected capture time.

The first fundamental step in order to design an optimal search strategy is to model the target motion. In fact, in many practical cases where a searcher is called to find a target, not only the position of the target is unknown but also its motion law. For this reason, a good approximate model is crucial to design a customized strategy which can exploit the available information.

In this paper, we consider different possible motion models and we study the effect that they have on the search strategy and on the expected capture time. In particular, the three different models for the target's motion we take into consideration are the following: 1) Stationary target; 2) Adversarial target; 3) Random moving target. All these models reflect general properties of the target and do not require a perfect knowledge of its characteristics. However, as we will show, they still contain important information for the searcher. These cases will be analyzed in detail in the following section.

This search problem has been inspired by our ongoing work on finding radio-tagged invasive fish in inland waters [20]. The task is carried out by using an Autonomous Surface Vehicle (ASV) equipped with antenna (see Fig. 1). In this application, no reliable motion model for the target (fish, in our case) is available. In fact, one of the final goals of this project is to provide tools for biologists to develop a motion model.

IV. Search Strategies

In this section we describe three different models for the target motion, which are analyzed separately. For each of them we then propose possible search strategies, discuss their properties and analyze expected capture time.

A. Stationary Target

A first possibility to model the target's motion is to assume that it is stationary. In our motivating problem, where the target is represented by a fish, this case is verified especially during the winter. Indeed, it has been observed that, during the winter months, these fish tend to aggregate and it is common to find them in large groups which remain stationary for many days [4]. If no information is available on the initial location of the target, an optimal search strategy is trivially represented by a deterministic precomputed path which covers the entire environment without revisiting any location. A classical example is the boustrophedon cellular decomposition [6]. For a square environment we can discretize Ω along the y-axis in $\hat{L}$ uniform rows, such that the distance between consecutive rows is equal to the sensing radius R_s. Then, the capture time will be at most $\hat{L}L$ and $\frac{\hat{L}L}{2}$ on average.

On the other hand, if the searcher has a non-uniform probability distribution for the target's initial location, different, more efficient strategies can be implemented and a wide literature is present on this subject (see Section II).

B. Adversarial Target

Another important model for the target's motion is to treat it as an adversary which actively tries to avoid the searcher. This could be possible assuming that the target is aware of the search strategy and it can plan the best possible strategy to escape capture by the searcher. Even though a perfectly adversarial target is unlikely in many real-world applications, this analysis remains very important because it can represent a worst case scenario. Indeed, a search strategy with capture (if it exists) guarantees the existence of a finite capture time for all possible trajectories of the target.

We now present a search strategy for capturing an adversarial target in the square environment Ω. If the target has a longer sensing range than the searcher it can easily avoid capture. Therefore, we assume that the players have the same sensing range R_s and so they do not observe each other until capture.

The optimal search strategy for this game is presented in [10]. We present here an adaptation of this strategy which yields a simple analysis in square environments. Let us consider the same discretization along the y-axis previously introduced. Therefore, we will have $\hat{L} = O(L)$ rows. The search strategy is divided into rounds. Each round has two parts which take exactly L steps. The searcher starts on the left boundary of the environment. At the beginning of each round, a row is randomly selected with a uniform distribution and the searcher travels to this row and waits until the end of the first part. In the second part, it sweeps the entire selected row. The searcher repeats this strategy until a capture occurs.

The capture probability of this strategy is computed as follows. We associate a "virtual searcher" with each random choice of row. At each round, only one virtual searcher is instantiated (i.e. it corresponds to the physical searcher). Hence, there are $\hat{L}-1$ virtual searchers. Every searcher occupies a different row and starts sweeping its row at the same time. In this way, the sensing areas of the all $\hat{L}-1$ virtual searchers and the real one create a continuous frontier which entirely sweeps the environment along the x axis. Regardless of the target's trajectory, at a given time the target will hit this frontier. As a result, one searcher among this virtual team will capture the target at every round. Note that the equal sensing range assumption implies that the target can not differentiate among virtual searchers choices and infer which virtual searcher is active.

The probability that such a capture happens on the row occupied by the real searcher is $1/\hat{L}$. The number of rounds before having a capture is geometrically distributed. Its expected value is equal to $\hat{L}$. Since each round lasts $2L$ steps, the expected capture time is

$$2L\hat{L} = O(A), \tag{2}$$

where A is the area of Ω. It is worth mentioning that no assumptions on the players' velocity were necessary in our analysis, and the same result holds also for a target which can move faster than the searcher.

This strategy is asymptotically optimal because by simply staying put at a randomly chosen location, the target can guarantee $O(A)$ capture time. A tighter analysis based on an improved evader strategy can be found in [10].

C. Random Motion Model

A third option to model the target's motion is a stochastic motion law. In many cases, even though the target's exact motion is unknown, it is still possible to approximate the motion of the target using probabilistic laws. For example, the movement of some animals can be approximated by stochastic models [13], [5].

In this paper we focus our attention on the random walk motion. Our aim is to show that even this partial probabilistic information about the target's motion and its expected behavior can be exploited to significantly improve the search strategies. In particular, we analyze two main ways to take advantage of the available information: employing a stationary searcher and designing adaptive strategies based on the belief of the target's position.

In order to proceed with the analysis of the stochastic motion model, we discretize the environment and describe the discrete process on a grid in the following way. Let the space Ω be represented by a grid $G = (V, E)$ composed of N nodes and consider a simple random walk on G. Let x_0 be the initial location of the random walk, which is completely unknown or follows a given non-uniform probability distribution. At the t-th time step let $1/d(x_t)$ be the probability to move from the current location x_t to one of the neighbors, where $d(i)$ is the degree of the node i. The sequence of random nodes $x_t, t = 1, 2, ...$ is a Markov chain. We then denote with P_t the distribution of x_t

$$P_t(i) = Prob(x_t = i)\,. \tag{3}$$

Denoting by $M = (p_{ij})_{i,j \in V}$ the matrix of transition probabilities for the Markov chain, we have:

$$p_{ij} = \begin{cases} \frac{1}{d(i)} & \text{if } ij \in E \\ 0 & \text{otherwise}\,. \end{cases} \tag{4}$$

Before we present the two strategies, an important result on this problem that we want to report is the lower bound on the expected capture time $\langle T \rangle$, which has been presented by Kehagias et al. in [14]. This result, which holds for a general graph composed of N nodes, with a maximum degree Δ, a minimum degree δ and such that $\frac{\Delta}{\delta(N-1)} \leq \frac{1}{24}$ states that:

$$\langle T \rangle \geq \frac{\delta(N-1)}{7e\Delta}\,. \tag{5}$$

For a square grid, this expression becomes:

$$\langle T \rangle \geq \frac{N-1}{14e}\,. \tag{6}$$

1) Stationary Searcher: A first important consequence of assuming the target as non-adversarial is that a viable strategy can be to just wait for it. In particular, if we consider valid the hypothesis that the target moves as a random walk in a bounded environment, we know that eventually it will visit every point in the search space with probability one. Two results already known in literature are particularly interesting for our analysis: given a random walk on a bounded two-dimensional grid composed of N nodes, *i)* defining the hitting time $h(i, j)$ as the expected time for a random walk starting at i to arrive to j for the first time, the maximum hitting time h_{max} over all the pairs (i, j) is $\Theta(N \log N)$; *ii)* the cover time, defined as the expected time to visit every node in the grid is $\Theta(N \log^2 N)$. See [15], [2] for more details.

Even though this expected time can seem very large compared with more efficient dynamic strategies, a stationary searcher can still be an interesting solution in some cases. One of them could be the case of a searcher with a limited energy budget: assuming that usually sensing without moving is significantly less costly than continuously moving, if the initial energy does not allow the searcher to explore an important part of the environment, a more efficient solution can be to employ the available amount of energy to sense in a fixed location. In Section V we will analyze in a quantitative way this aspect.

Furthermore, this kind of knowledge about the target model, can also allow employing one, or more, static sensors distributed over the environment, instead of a more expensive mobile robot.

2) Adaptive Searcher: An alternative way to exploit the probabilistic information about the target is to update a probabilistic belief function of the target's position during the search and to use it to construct adaptive strategies. In this way it is also easy to incorporate possible a priori information about the initial position of the target.

The strategy we propose is to update the belief over the entire grid. For each node, we add the probabilities of its neighbors weighted by the transition probability given in eq. (4). We then move toward the global maximum of this function. The drawback of this kind of strategy is that it can be very computationally expensive for large environments. This drawback can be mitigated by computing the search path in advance for a given initial distribution and search time.

We have also considered a simple greedy strategy, where at each step the best node among the neighbors is selected. We observed that it is very sensitive to local minimum problems and get stuck (e.g. along the boundary) easily. As a result, we have decided that it is not a promising search strategy.

It is important to mention that these strategies are not the solutions of the optimization problem of minimizing the expected capture time (or related problems such as maximizing the capture probability in a given time), which is an open problem.

V. Simulation Results

In this section we focus our attention on the random moving target and we study in simulation the strategies discussed in Section IV. Specifically, the considered strategies are: 1) Simple random walk; 2) Boustrophedon path; 3) Adversarial-based path; 4) Random direction; 5) Belief-based path.

Our intent is to investigate the behavior of the mean capture time for these strategies as a function of the environment size. For our study we consider a square grid composed of N nodes, on which searcher and target can move. The target's initial location is fixed randomly, assuming a uniform probability distribution over the N nodes. The searcher starts from the point $s_0 = (1, 1)$. The sensing radius R_s is fixed to 1, which means that the searcher can also sense its four neighbors. Searcher and target move at the same velocity, which is assumed to be unitary.

In the first strategy the searcher moves in the same way as the target, with the same probabilities at each time step. The adversarial-based path is the strategy described in Section IV-B. In the random direction, the searcher picks a random, uniformly-distributed point on the boundary, it reaches it and repeats this process. In the belief-based strategy, at each iteration the maximum of the belief is identified and the searcher goes toward its direction. Fig. 2 shows the simulation results, where the numerical values for each strategy have been obtained as an average over 10^4 trials.

It is not surprising that the best performance is achieved by using the belief function to design an adaptive strategy. However, it is interesting to note that its performance is still close to the one corresponding to randomized strategies, i.e. adversarial-based and random direction. In particular, on average, the adversarial-based strategy is able to catch the target in a time 20% higher than the belief-based one. It is worth noting that these two non-adaptive random-based strategies, for this case with a randomly moving target,

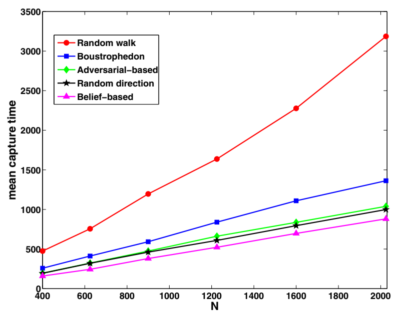

Fig. 2. Behavior of the mean capture time as a function of the number of nodes N. This plot shows five different searcher strategies: moving as a random walk, boustrophedon path, the random strategy presented in Section IV-B and the adaptive strategy based on the belief function. The sensing radius R_s is fixed to 1.

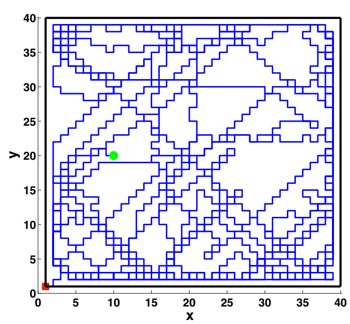

Fig. 3. Belief-based search path: at each time step, the searcher moves toward the global maximum of the updated belief of the target's location. The red square represents the initial position, the green circle the final one.

achieve a better result than the boustrophedon path. This phenomenon can be explained as follows: when the searcher sweeps an entire row without finding the target, the value of the belief drops also on the adjacent rows. Thus, instead of passing to sweep the following row, a higher capture probability can be obtained exploring more distant regions.

Another important remark is that the last four strategies show a mean capture time linear in the number of nodes N, but searching as a random walk requires a time more than linear. This fact can be easily explained because in this case the global system, searcher plus target, is still a random walk, with a different diffusive law. As a result, an analysis similar to the one presented for a stationary searcher can be applied and even if the expected capture time for this case is lower, the behavior is always more than linear in N.

In Fig. 3 is shown an example of a belief-based search path in a square environment of side $\sqrt{N} = 40$. It is possible to see that the searcher always avoids the boundaries of the environment. This is due to the lower degree of the boundary nodes, and so also the lower probability to find the target.

An aspect that is worth considering is the robustness of this belief-based strategy with respect to changes in the expected target motion, i.e. when the random walk probabilities are different from the ones used to update the belief. This situation reflects a very realistic scenario, where the target does not behave exactly as the model adopted but only close to it. The results of this analysis, now in terms of capture probabilities as a function of time, for a searcher moving in a square of side $L = 50$, are reported in Fig. 4. The belief-based strategy is obtained considered the simple random walk as described in eq. 4, but now the target can also have a non-zero probability of maintaining is position using the transition probabilities:

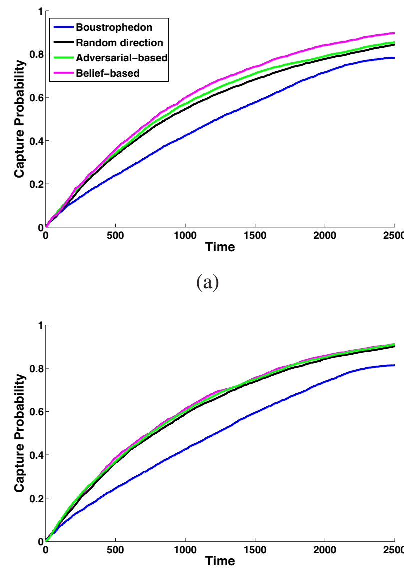

(a)

(b)

Fig. 4. Behavior of the capture probability in function of time with $w = 0$ (a) and 0.2 (b).

$$p_{ij} = \begin{cases} w & \text{if } i = j \in E \\ \frac{1-w}{d(i)} & \text{if } ij \in E \\ 0 & \text{otherwise}, \end{cases} \tag{7}$$

where $w = 0, 0.2$ in Fig. 4(a), (b) respectively. *These results show that, even though the belief-based strategy achieves the best result, starting from a uniform probability distribution for the initial target's location for a perfect simple random walk motion, the gain with respect to randomized strategies is lost as soon as a slight modification in the transition probabilities is present.*

Finally, we want also to compare in simulation the performance of a stationary searcher with respect to a moving strategy. As stated in the analysis presented in the previous section, the expected time for the target to hit a stationary searcher is $\Theta(N \log N)$. On the other hand, for a searcher adopting a moving strategy such as the adversarial-based or the belief-based, we expect that the time drops to $\Theta(N)$. This was proven for the adversarial-based strategy in Section IV-B and confirmed for both strategies in simulation (Fig. 2). To verify this different behavior as a function of the number of nodes N between a stationary searcher and the belief-based strategy, we consider the ratio:

$$\frac{T_{\text{stationary}}}{T_{\text{belief}} \log N}, \tag{8}$$

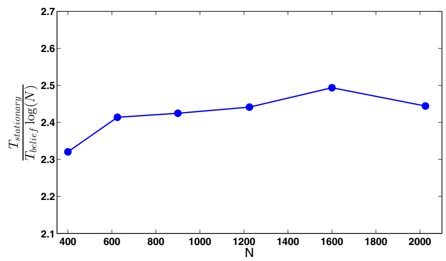

Fig. 5. Behavior of $\frac{T_{\text{stationary}}}{T_{\text{belief}} \log N}$ as a function of N, where $T_{\text{stationary}}$ and T_{belief} represent the mean capture time for stationary searcher and belief-based strategy respectively.

where T_x represents the mean capture time for the x strategy. We expect this ratio to be a constant value K as a function of N. The result is shown in Fig. 5. It is possible to see that the attended behavior is verified and in particular, for our case, we have $K \sim 2.4$. This result is particularly interesting, not only to confirm the theoretical behavior, but also to obtain a quantitative comparison between the two strategies. As mentioned in the previous section, this could be important for a limited-energy search mission. For instance, in this case, if the cost of completing a move action c_m compared with the cost of sensing keeping the position c_s verifies

$$c_m \gg c_s K \log(N) \tag{9}$$

from an energetic point of view, it would be more efficient to wait for the target instead of moving to search for it. However, it is important to note that this result can strongly depend on the real target's motion. If it does not present all the characteristics of a random walk (for example if a strong bias is present) the eventual capture could not be assured.

VI. Conclusion

In this paper we studied how the choice of target model affects the strategy of a mobile searcher which tries to minimize the time of capture. We considered three target models: stationary target, adversarial target and target moving as a random walk. The objective was to provide a review of existing results in the field, adapting some of them to our particular case, and present possible strategies for the target models. In the case of a random-walking target, we also present a comparison in simulation of viable search paths. Our results show that when searching for a random-walker, an effective strategy is to maintain a distribution of the target locations and to move toward the maximum of this distribution. However, this strategy is sensitive to errors in the target motion model. Randomized (adversarial) strategies achieve a performance close to that of a belief-based strategy. Since they are agnostic to the target model, they provide viable general search strategies.

Finally we proposed an analysis of this case also from an energy point of view, considering the case of a stationary target and comparing it with a moving strategy.

As future work, we aim to further investigate the problem of a randomly moving target from an energy point of view. Assuming a limited energy budget for the searcher and different costs for stay and moving actions, as in [19], our intent is to try to design efficient strategies which combine the exploration and stationary modes in order to increase the capture probability. Another possible research line is to study the effects of a different stochastic motion law for the target. An interesting alternative to the random walk model is that represented by Levy flights.

References

[1] M. Adler, H. Racke, N. Sivadasan, C. Sohler, and B. Vocking. Randomized pursuit-evasion in graphs. *Combinatorics Probability and Computing*, 12(3):225–244, 2003.

[2] N. Alon, C. Avin, M. Koucky, G. Kozma, Z. Lotker, and M. R. Tuttle. Many random walks are faster than one. In *Proceedings of the twentieth annual symposium on Parallelism in algorithms and architectures*, pages 119–128. ACM, 2008.

[3] J. Anlauf. Asymptotically exact solution of the one-dimensional trapping problem. *Physical review letters*, 52(21):1845–1848, 1984.

[4] P. G. Bajer and P. W. Sorensen. Recruitment and abundance of an invasive fish, the common carp, is driven by its propensity to invade and reproduce in basins that experience winter-time hypoxia in interconnected lakes. *Biological Invasions*, 12(5):1101–1112, 2009.

[5] F. Bartumeus, M. G. E. da Luz, G. Viswanathan, and J. Catalan. Animal search strategies: a quantitative random-walk analysis. *Ecology*, 86(11):3078–3087, 2005.

[6] H. Choset. Coverage of known spaces: the boustrophedon cellular decomposition. *Autonomous Robots*, 9(3):247–253, 2000.

[7] T. Chung, G. Hollinger, and V. Isler. Search and pursuit-evasion in mobile robotics. *Autonomous Robots*, 31(4):299–316, 2011.

[8] T. H. Chung and J. W. Burdick. Analysis of search decision making using probabilistic search strategies. *IEEE Transactions on Robotics*, 28(1):132–144, 2012.

[9] A. Gabel, S. Majumdar, N. Panduranga, and S. Redner. Can a lamb reach a haven before being eaten by diffusing lions? *Journal of Statistical Mechanics: Theory and Experiment*, 2012(05):P05011, 2012.

[10] S. Gal. Search games with mobile and immobile hider. *SIAM Journal on Control and Optimization*, 17(1):99–122, 1979.

[11] G. Hollinger, S. Singh, J. Djugash, and A. Kehagias. Efficient multi-robot search for a moving target. *The International Journal of Robotics Research*, 28(2):201–219, 2009.

[12] R. Isaacs. *Differential Games: A Mathematical Theory with Applications to Welfare and Pursuit, Control and Optimization*. John Wiley and Sons, 1965.

[13] P. Kareiva and N. Shigesada. Analyzing insect movement as a correlated random walk. *Oecologia*, 56(2-3):234–238, 1983.

[14] A. Kehagias, D. Mitsche, and P. Pralat. Cops and invisible robbers: The cost of drunkenness. *Theoretical Computer Science*, pages 100–120, 2013.

[15] L. Lovász. Random walks on graphs: A survey. *Combinatorics, Paul Erdos is eighty*, 2(1):1–46, 1993.

[16] N. Noori, A. Renzaglia, and V. Isler. Searching for a one-dimensional random walker: Deterministic strategies with a time budget when crossing is allowed. In *International Conference on Robots and Systems (IROS)*, 2013.

[17] S. Redner. *A Guide to First-Passage Processes*. University Press, Cambridge, 2001.

[18] S. Redner and P. Krapivsky. Capture of the lamb: Diffusing predators seeking a diffusing prey. *American Journal of Physics*, 67:1277–1283, 1999.

[19] A. Renzaglia, N. Noori, and V. Isler. Searching for a one-dimensional random walker: Randomized strategy with energy budget. In *IEEE Conference on Intelligent Robots and Systems (IROS)*, 2013.

[20] P. Tokekar, E. Branson, J. Vander Hook, and V. Isler. Coverage and active localization for monitoring invasive fish with an autonomous boat. *IEEE Robotics and Automation Magazine*, 2012.

[21] G. H. Weiss, S. Havlin, and A. Bunde. On the survival probability of a random walk in a finite lattice with a single trap. *Journal of statistical physics*, 40(1-2):191–199, 1985.