

March 5, 2010

To the Workshop Organizers or Whom it May Concern,

We are looking forward to attending your workshop to discuss our recent work on probabilistic pursuit of adversarial evaders. Attached we have included some of our work which is currently under review at RSS. Our presentation at the workshop would borrow heavily from the RSS paper, but would also include some additional work on dynamically responding to evaders of varying adversarialness. Please let us know if you have any questions or concerns.

Sincerely,

Johannes Strom, Ryan Morton, Kevin Reilly
Professor Edwin Olson
Department of Computer Science and Engineering
University of Michigan

`{jhstrom,rmorton,keegreil,ebolson}@umich.edu`

Online Probabilistic Pursuit of Adversarial Evaders

Blind Submission #10

Abstract—The pursuit-evasion problem consists of a team of pursuers maneuvering to capture one or more evaders. Solutions to the pursuit-evasion problem help search and rescue teams quickly find survivors, sentries protect against intruders, and law enforcement apprehend suspects. Early work considered the case of adversarial evaders, though these approaches become prohibitively expensive when scaling to large environments and/or multiple pursuers. More recently, probabilistic pursuit methods have yielded computationally efficient approaches, but have sacrificed the ability to deal with adversarial evaders.

We describe a new probabilistic method that solves the adversarial pursuit-evasion problem without the need for costly adversarial searches (e.g., minimax). In particular, we show how a random-walk model of evaders can be transformed into a model that reflects the behavior of an adversarial agent. Our approach allows pursuers to characterize their opponents’ cunning, allowing faster capture times in the case of sub-optimal evaders. Finally, we demonstrate our algorithm’s performance on a number of benchmark problems, and verify our implementation on a robot.

I. INTRODUCTION

In the pursuit-evasion problem, a team of agents (the pursuers) are collectively trying to “capture” another set of agents (the evaders). In addition to the many military applications, pursuit-evasion can be applied to search and rescue, surveillance and monitoring, and other related problems. However, computing good plans for a team of coordinating pursuers can be computationally prohibitive, particularly if the evaders are adversarial.

There are many variations of the pursuit-evasion game. For example, in search variants of the problem, the location of the evaders must be discovered; in pure pursuit variations, the position of each agent is known by all the other agents. The speed of the evaders also varies; some formulations limit the speed of the evaders while others assume the evaders can move at infinite speed. Perhaps the most critical question relates to the behavior of the evaders: are they moving about the world randomly, patrolling the world in a fixed pattern, or are they explicitly trying to avoid capture?

In this paper, we focus on the pursuit-evasion game that is central to the upcoming MAGIC 2010 robotics competition [1]. This game combines elements in a novel way that poses challenges for existing algorithms, yet, at the same time, is motivated by a real-world application. For example, the layout of the world and the real-time position of all the evaders is known to all agents due to a reconnaissance aircraft. The evaders are not moving randomly, but are executing a motion plan that is unknown to the pursuers.

This variation poses challenges for existing algorithms. Visibility-based approaches [2] provide guaranteed capture, but do not scale well to large numbers of pursuers or evaders. Further, these approaches cannot handle cases where guaranteed capture is impossible, due either to the structure of the




Fig. 1. (a) Previous probabilistic pursuit methods assumed random evaders. (b) We show how to extend this model to adversarial evaders which require more teamwork to catch.

world or to an inadequate number of pursuers. Probabilistic methods provide a way forward, but they typically assume a known (or approximately known) *random* motion model for the evaders. As we will show in this paper, an evader that moves randomly is generally very easy to catch (each “bad” move allows the intelligent pursuers to close in). Unfortunately, the formulation of existing probabilistic methods do not easily allow modeling adversarial motion models.

We extend recent graph-based works by Vieira and Hollinger [3],[4] and propose a new model of adversarial evaders which enables efficient computation of a pursuit strategy that scales linearly with the number of pursuers or evaders. Our probabilistic model of adversarial evaders extends the efficient, though non-optimal, random diffusion model discussed in detail by Hollinger, while maintaining the same bounds on complexity. Vieira’s work explicitly handles adversarial evaders whose locations are known, but in a way which requires a large ratio of pursuers to evaders. The method we propose efficiently enables the capture of an arbitrary number of adversarial evaders with minimal pursuers. Furthermore, our proposed model allows incorporating a priori information about the adversarialness of an opponent to minimize capture time. On the one hand, Hollinger points out that overestimating the skill of an evader results in overly conservative pursuer strategies that take many moves to carry out [5]. On the other hand, underestimating an evader’s motion model as random is a poor approximation for *most* real world agents and leads to longer capture times in such situations. To perform well in both scenarios, Hollinger employs a suboptimal division of the searching robots into two sub-teams, the first of which greedily pursues highly-probable locations while the second systematically forces capture. Our proposed method does not require a partitioning of the robots: all robots run the same algorithm and the predicted skill of the evaders is quantified and formally integrated into the rewards planning.

In summary, the key contributions of this paper are:

- We describe a novel model of an adversarial agent with variable skill. Given a single skill parameter, the behavior

of the adversarial evader spans from purely random to horizon-optimal.

- We show how this adversarial model can be integrated into a rewards-based planner in a principled fashion.
- We present a rewards-based planning algorithm that efficiently computes plans for pursuers, even against adversarial opponents.

The remainder of this paper follows the following format: a discussion of previous work is presented in Section II, followed by a detailed description of the planning algorithm and the adversarial evader model in Section III. We demonstrate the performance of our method with results in both simulation and on a robot platform in Sections V and VI, respectively. Finally we place our method in context and discuss future work in Section VII.

II. PRIOR WORK

The domain of pursuit-evasion games covers a large body of work with diverse applications. Sometimes referred to as the “Museum Gallery” problem [6], or the “Cops and Robbers” problem, these problems generally consider a group of one or more pursuers who are trying to capture or locate a group of one or more evaders. Some formulations involve an explicit search component to find the evaders, while others use sensor nets or limited sensor measurements to inform pursuit [7, 8, 4]. Most formulations use a known map, but the nature of the environment differs; Parsons’ seminal paper formulated the problem on a graph, while LaValle explored the problem in simple geometric environments [8, 9, 2]. Isler further examined the geometric case when communication between pursuers is limited [10], while Parker extended the same environment using probabilistic methods and gradient descent to the tracking of multiple robots [11]. Other distinctions between approaches include probabilistic vs. exhaustive search methods, such as those recently explored by Hollinger or Vieira, respectively [5, 3].

Most applicable to this paper are the contemporary formulations of the pursuit-evasion-on-a-graph problem in probabilistic terms, such as in the greedy approaches discussed in [12, 13]. Discretizing continuous real world environments into graphs provides a crucial simplification which allows efficient computation of pursuit strategies [14]. Vieira’s recent work on graph-based pursuit is also extremely relevant as it addresses capture of multiple optimally adversarial evaders [3]. Probabilistic formulations, such as in Hollinger’s extensive work, allow online computation of effective, but sub-optimal pursuit strategies against randomly moving evaders [4, 5, 15, 16]. These probabilistic methods have many advantages, among which are lower time complexity and applications to highly stochastic and cluttered environments [16].

Hollinger’s work exhibits both search and pursuit variants of the problem. In the former, agents are searching a graph for the location of an evader whose location is unknown and whose motion is assumed to be random [16]. In the pursuit variant, pursuers chase evaders whose position is estimated with range measurements [4]. In both variations, each pursuer plans a trajectory to maximize the collective rewards of the team.

The location of the evader is represented as a probabilistic *belief* over the nodes of the graph. Conveniently, the *belief* representation applies to both variants. Over time, this *belief* is propagated by the given motion model. Rewards for candidate trajectories are calculated by evaluating the probability of capturing the evader given the belief of the evader’s location, and a candidate motion plan. Hollinger also applies a discount factor γ to future rewards, an approach which was shown by [17] to make searches in large or infinite state spaces tractable. Collaboration between pursuers is *implicit* [5]: only the most recently computed plans of other agents are considered when planning a single agent’s future path. As a result, optimality is lost, but the total computation time grows only linearly with the number of pursuers. A generalization of this coordination algorithm was also successfully applied to multi-agent POMDPs [18]. Alternatively, explicit collaboration can compute optimal plans. However, explicit collaboration has a branching factor proportional to the number of agents, and consequently has exponential runtime with respect to the number of pursuers.

Hollinger’s approach works well when the target is assumed to behave in a non-adversarial manner. However, if the agent is even slightly adversarial, the random diffusion model is no longer valid and average-case capture is not guaranteed, since a nominally smart evader will likely escape much of the time. To guarantee capture, Hollinger extends the probabilistic search method with a guaranteed search algorithm that partitions the graph into trees which can be decontaminated systematically. This yields bounded capture time for any type of agent, but carries no guarantee for optimality. Furthermore, when agents are only partly adversarial, this method generates search strategies which take considerably longer than necessary. Although Hollinger does not explicitly provide a method for handling multiple evaders, the specified search strategy can be computed a priori, and as such would apply to multiple evaders seamlessly. Hollinger’s work has been extended to the case where two robots are required to capture a potentially “adversarial” evader [19]. However, [19] actually deals exclusively with the same randomly moving evaders discussed above.

In contrast to Hollinger, Vieira only studies the pursuit variant of the problem where the locations of the evaders are assumed to be given exactly from a sensor network [3]. Vieira demonstrates a non-probabilistic strategy which explicitly handles multiple evaders by computing optimal, shortest time capture trajectories for a small number of pursuers. Computation of optimal paths is exponential in the number of pursuers and evaders. To scale to multiple evaders, Vieira partitions the problem by assigning a small group of pursuers to each evader. While the partitioning method provides a method to keep the computation time feasible, for each additional evader, this solution requires at least $c(G)$ additional pursuer robots¹, resulting in significantly sub-optimal performance when there are many evaders. Furthermore, as the number of evaders grows, the requisite number of pursuer robots required by that algorithm can quickly become impractical. In domains where

¹The capture number $c(G)$ for a graph is the smallest number of pursuers required to guarantee capture of an evader on a graph G .

the capture number grows large, the number of pursuer robots assigned to each evader will quickly grow beyond the available computational capacity.

In summary, there has been a great deal of work done on the pursuit evasion problem. This paper extends probabilistic pursuit methods to multiple, adversarial evaders, which previous authors have identified as an important problem [5, 16]. Previously, these topics have been studied individually [3, 16], but have not yet been unified.

III. PROBLEM FORMULATION

Building on the pursuit strategy pioneered by Hollinger, we formulate the problem in terms of a rewards-maximizing search [16]. In other words, the strategy systematically evaluates all possible plans to find the one with the highest reward. Rewards are associated with the capture of an evader (R_{evader}): we additionally incorporate a discount factor $\gamma \in [0, 1]$ such that any immediate rewards earned at time step t are discounted by a factor of γ^t . Discount factors are particularly applicable in robotics problems: failures and unexpected delays can easily derail a plan, making it possible that a reward might not actually be earned. Furthermore, they encourage agents to develop plans with faster capture times.

As in [5, 14], we assume that the physical world can be abstracted to a graph, where nodes represent locations in the world and agents can move between nodes when they are connected by an edge. This assumption can be motivated by the fact that a real world environment can be decomposed into convex regions, each of which represents a node. Evaders are considered “caught” when a pursuer occupies the same node as an evader. In a building, for example, rooms and hallways can constitute such convex regions.

The problem for each of the pursuers is to choose a sequence of moves that maximizes the expected reward of the entire team of p pursuers. If each pursuer has a available actions, the branching factor of the search is $b = a^p$. As a result, directly searching the space of joint pursuer trajectories is infeasible. Classical adversarial search would also branch on the actions of each evader, giving a depth of $b = (ae)^p$ for e evaders. The complexity is even worse if we allow the evaders to be jointly adversarial, since that increases the branching factor even further. However, in this paper we focus on the former case, where the evaders are not jointly adversarial.

Our approach addresses these complexity issues in two ways. First, to mitigate the exponential cost associated with many pursuers, we use the same *sequential* planning technique as Hollinger [15]: each pursuer computes a plan for themselves given fixed plans for its teammates. The pursuers then exchange plans, move to a new node, and plan once again. A similar iterative approach is analyzed in detail by [18]. In the non-adversarial case, this distributed approach reduces the branching factor of each search to just a , though there is naturally no guarantee of optimality. Second, to address the high branching factor incurred by adversarial evaders, we will show later how to model the actions of the evader such that the branching factor of the search tree remains $b = a$.

A. ϵ -Optimal Planning

The depth of the search also plays a critical role in the complexity of the search. Rather than have the user specify a search depth, our approach is to have the user bound the regret of the computed plan. In other words, we will guarantee that the reward of the computed plan is within ϵ of the optimal plan. Given a mechanism for computing an upper-bound on the maximum reward that can be achieved from a state, we can immediately compute the necessary search depth required.




Fig. 2. Discounted rewards as depth increases. The immediate reward is discounted by γ^d and the total reward for a plan is the sum of incremental rewards along the path. The value of $R_{i,t}$ corresponds to the cumulative reward at node i at time t . If at any node the total possible reward for continuing down this path is less than ϵ the rest of that path will be pruned.

Suppose that the maximum value of any node (i.e., its immediate rewards plus discounted future rewards) is V and that the discount factor is γ . The discounted reward at depth d is just $\gamma^d V$; once this is less than ϵ , we have bounded our regret by ϵ . This leads to a maximum search depth of:

$$d_{max} = \log_{\gamma} \frac{\epsilon}{V} \quad (1)$$

We can use discounted rewards not only to bound the maximum search depth, but also to prune search nodes when they could be no more than ϵ better than the best plan known so far. This yields a depth-first search with a strong pruning criterion.

B. Modeling Evaders

As we mentioned before, optimal handling of adversarial evaders requires branching on each possible opponent action, which renders a search to a useful depth infeasible. Previous probabilistic methods side step this issue by assuming non-adversarial evaders [15]. These works consider a known diffusion model for a randomly moving evader: at each depth, the evader’s action consists of a “diffusion” to adjacent nodes. Assuming the transition probabilities for the evaders are known (e.g. random), this leads to an exact *belief* representation of the evader’s possible locations. Furthermore, such a representation allows the branching factor of the search to only depend on the number of actions available to the pursuer who is formulating a plan [5]. However, *most* real world agents are significantly

harder to catch than a randomly moving opponent. Even an agent following a set of scripted actions can be substantially harder to catch than an agent that is performing a random walk. In the extreme case, if the evaders are actively trying to avoid capture, the method falls apart completely.

The goal of this work is to demonstrate how to handle adversarial evaders using this same framework, requiring only a branching factor of $b = a$.

1) *Adversarial Evaders*: In the probabilistic graph search framework in [5], the pursuers compute the belief state over the position of each evader on the graph at each node in the search tree – assuming that the evaders move according to a fixed set of transition probabilities (e.g. random). Given that the graph has n nodes, the belief state has length n . We now extend this representation so that the belief state has length $n + 1$ to include the cumulative probability that the evader has been captured, $P(\text{capture})$. In the case of the random dispersion model, this is straightforward to determine exactly. As agents explore the tree of possible actions, the belief is updated according to the evader’s motion model. For example, the belief corresponding to the evader’s presence in a particular location on the graph is dispersed to its neighbors according to the random model. Given this model, we track the exact probability of the evader being in that location, or each graph node. However, during the propagation of the belief, the portions which intersect the planned path of the pursuing team are considered captured, and added to P_{cr} . P_{cr} is the probability that the evader is *captured*, given the *random* dispersion model. This is in contrast to $P(\text{capture})$ which is the *true* probability of capture. When the motion model is known to be random, P_{cr} is exactly $P(\text{capture})$. Unfortunately, in the case of adversarial evaders, P_{cr} is necessarily an overestimate of $P(\text{capture})$. This is because such agents are likely to actively avoid actions which will result in capture.

To properly handle adversarial evaders, we need to estimate $P(\text{capture})$ correctly. However, exactly determining $P(\text{capture})$ requires access to– and the ability to run– the exact algorithm that each evader is running. Both of these assumptions are not guaranteed, since the evader could be running any unknown algorithm, particularly one which is computationally intractable on the pursuers hardware (e.g. an adversarial search with a high branching factor).

To calculate $P(\text{capture})$ in the general case we introduce a different evader model which incorporates agents of arbitrary adversarialness using a single skill parameter, S . Using this model, the pursuing team can accurately *model* the actions of the evaders assuming that S is known *a priori*. This assumption is often reasonable because in many domains, prior information about the adversarialness of opponents is available. For domains where this is not the case, we demonstrate an experimental method for determining S for a particular adversary in Section V.

To define our model explicitly, consider that the evaders in our model build a tree of their possible actions, but that this tree is limited to S leaves (e.g. due to computation limitations). The actions in this tree are selected randomly. However, the agent picks the best of those sampled actions at each turn. Since each of the actions is a random trajectory, the

probability that one action leads to capture is precisely the P_{cr} computed by the random diffusion model. Our model assumes that each evader will pick an evading strategy, if one exists. The probability that *all* of the S trajectories lead to capture is shown in Equation 2.

$$P(\text{capture}) = (P_{cr})^S \quad (2)$$

Thus, given our evader model, we can correctly estimate $P(\text{capture})$. Although the model may not correspond exactly to the evader’s actual algorithm, a key point is that it necessarily encapsulates all possible adversarial evader models. This includes acting randomly, when $S = 1$. As S grows, increasingly adversarial evaders can be modeled. At the limit, when $S \rightarrow \infty$, the evader can examine all possible trajectories, corresponding to a perfect adversary. To see the effect on $P(\text{capture})$, consider the behavior of Equation 2 as a function of P_{cr} and S (see Figure 3): Suppose that P_{cr} is less than 1.0 (i.e., that there exists a trajectory that escapes). As S approaches infinity, the true probability of capture goes to zero.




Fig. 3. The effect of the adversarial skill parameter S on expected probability of capture, given the probability of capture P_{cr} from the random diffusion model in [4]. (See Equation 2.)

Using Equation 2 together with the search framework we outlined above, we are able to calculate the *cumulative* $P(\text{capture})$ for any evader at any node in the pursuer’s search tree. To compute the rewards associated with the candidate plan, we must sum the marginal rewards at each depth (see Figure 2). The *marginal* rewards gained at each depth i depend on the *cumulative* $P(\text{capture})_i$. That is, the marginal rewards are proportional to the *difference* between the cumulative capture probability at the current depth minus the cumulative capture probability at the last depth, and subject to the appropriate discount factor. For reward associated with the potential capture of a single evader this is:

$$R_{total} = R_{evader} \sum_{i=0}^d [P(\text{capture})_i - P(\text{capture})_{i-1}] \gamma^i \quad (3)$$

Where, $P(\text{capture})$ is determined using Equation 2.

Consider a pursuer chasing a possibly adversarial evader and must choose between three possible trajectories: A, B and C (see Fig. 4). Trajectory B initially has a low probability




Fig. 4. Sample trajectories and the associated rewards. A pursuer is choosing among three candidate sequences of action. Each node is labeled with an estimate of P_{cr} , the probability that, for that depth, the evader will be captured given the random motion assumption. The cumulative reward for each trajectory is shown both with and without the assumption that the evaders are adversarial. $S = 20$, $\gamma = .8$

of capture, P_{cr} , but eventually offers the best chance for capturing the evader. An evader acting under our model would need to have sampled one of the 10% of evading trajectories in order to escape. When S is large, B is clearly the best pursuit trajectory. On the other hand, if the evader is assumed to be moving randomly the chance of the evader choosing an escaping trajectory is relatively slim. The discount factor ensures that our pursuers will prefer trajectory C which offers a greater chance of capturing the random evader sooner than trajectory A.

C. Complexity Bounds

Our model offers a tremendous improvement in worst case bounds over previous adversarial search methods such as minimax: the worst case computation bounds of our algorithm is $O(b^d)$, rather than $O(b^{ed})$. This improvement is made possible because we employ implicit collaboration between the pursuers and because of our evader model. That is, we eliminate the need to branch on any of the evaders' actions, or on the actions of any the pursuer's teammates. As a result, given equal computational resources, our pursuers are able to search significantly deeper than traditional adversarial search algorithms. The reason we are able to improve these bounds is because we do not keep track of individual evader trajectories. Computing P_{cr} throughout the search process only encodes the information that the evader *could* escape, not exactly which path it must take to do so. Empirically, we do notice an effect on the runtime due to increasing S . This is because the pruning criterion discussed in Section III-A has the greatest effect when good plans are found near the top of the search tree. When faced with an adversary who the pursuers believe is highly cunning (S is large), highly rewarding plans can only be found by ensuring capture. When possible, such forced capture occurs only sparsely in the search tree, and often requires complex and lengthy plans to achieve. This reduces the effectiveness of the pruning criterion.

D. Multiple Evaders

The previously described single evader methods extend seamlessly to multiple evaders. As mentioned previously, our approach can easily incorporate belief states for multiple evaders. The most interesting aspects of our algorithm are manifested in cases with multiple evaders. Pursuers who are assuming random evaders will greedily pursue locations of high probability, regardless of which evader is expected to be there. Using our adversarial model, pursuers who are modeling the evaders with $S > 1$ will increasingly ignore nearby evaders who they predict they cannot catch on their own in favor of traveling larger distances to help another pursuer corner an adversary. Figure 1 shows an example of this in a simple environment from [4]. When the targets are assumed to be adversarial, both pursuers team up to "force" capture. As S grows, the pursuers will increasingly ignore strategies with low probability of capturing many evaders, in favor of strategies which guarantee capture of fewer evaders.

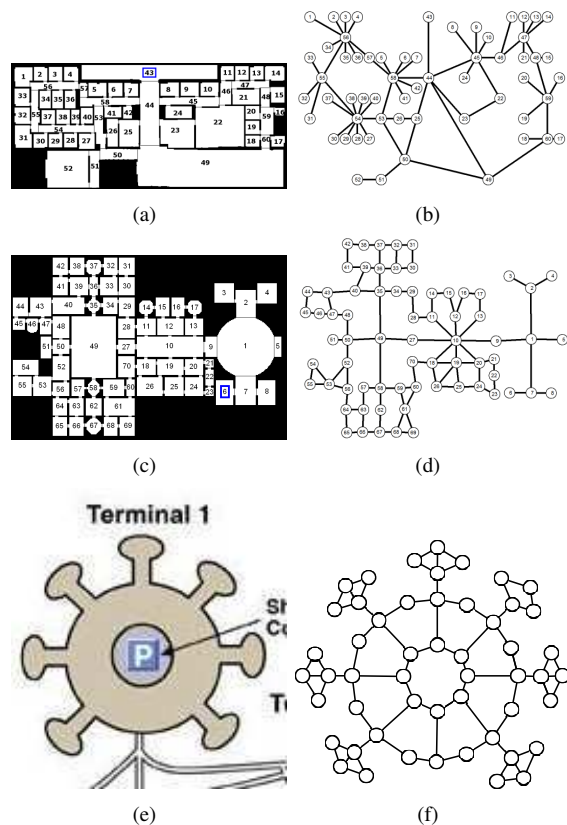


Fig. 5. Map and graph pairs used in this project. The graphs of Newell-Simon Hall at Carnegie Mellon University (a) and the US National Gallery Museum (c), were introduced by Hollinger and offer a good benchmark since these have been used in previous work. We also introduce a more uniformly complex map of the Charles de Gaulle airport in Paris (e).

IV. EXPERIMENTAL SETUP

We tested our algorithm both in simulation and on real robots. In simulation, we ran trials on the various standard graphs published previously by Hollinger. In addition, we




Fig. 6. Comparative run times (log-scale) of the pursuers’ algorithm (blue) and the evaders’ minimax algorithm (red) to varying depths. Note that our algorithm searches to a depth of 6 timesteps in the same time that minimax searches to a depth of 3.

introduce a new map of the Charles de Gaulle airport (CDG), whose complexity is more suitable for benchmarking adversarial evaders (see Figure 5). In each test, we simulated a pursuit evasion game between several evaders and several pursuers. For all the tests, we set a discount factor of $\gamma = 2/3$ and a maximum regret $\epsilon = .05R_{evader}$. By Equation 1, this gives a maximum search depth of 6, which remains tractable to simulate online on a 2.6 GHz Intel Core2Duo laptop.

We tested our pursuit algorithm against three different types of simulated evaders:

- 1) A minimax adversarial evader.
- 2) A “sampling” adversarial evader which has access to the pursuit algorithm described herein, and can sample T trajectories and pick an evading path. (Similar to the model described in Section III-B).
- 3) A fixed patrol (in the robot case).

For our tests, we allowed the minimax algorithm to run to depth 4, the maximum depth which allowed reasonable run-time. Given similar time constraints, the pursuers’ algorithms were able to search to depth 6 (see Figure 6). Here, *depth* refers to the number of future timesteps an agent can search. The “sampling” adversary was made very adversarial by having the capability to sample 100 possible trajectories. In all tests, a maximum time limit is set such that the evaders who are still alive at the time limit “escape”.

To benchmark our method, we tested our algorithm in a variety of simulated scenarios; for comparison with previous work we show results from both the search and pure pursuit variants of the problem. Hollinger’s recent work deals with the search variant of the problem, while Vieira and Hollinger’s early work deal primarily with the pursuit problem (where information about the location of the evaders is known)[4, 3, 16]. With regard to the determination of the skill parameter S , we demonstrate how to approximate S for a minimax agent of depth one. Finally, pertaining to the upcoming MAGIC competition, we also show a pursuit trial on real robots where human evaders execute fixed patrols.

For each simulated experiment, we conducted many random




Fig. 7. On the museum graph 5b, capture time decreases significantly as the pursuers’ model of the evader more accurately reflects their adversarialness. $S = 1$ corresponds to the random evader considered in Hollinger’s work, while higher S correspond to explicitly modeling the evaders as adversarial. In this trial, two pursuers team up to capture five independently adversarial evaders. Error bars reflect one standard error of the mean.




Fig. 8. For comparison, we validate our model by running 2 pursuers on the museum graph against 5 evaders running a fixed depth mini-max search with alpha-beta pruning. For this trial, the evader depth was limited to 4 for computational reasons. This graph shows matching trends to Figure 7, and provides a baseline for comparison to our 100-sample agent.

trials to determine the behavior of our algorithm in various situations. At the start of each trial, random, non-overlapping starting locations were assigned to pursuers and evaders. As in Hollinger’s work, the pursuers all start in the same location, while the evaders were dispersed throughout the remainder of the graph. The algorithm was then run for each value of the independent variable, given those starting locations. New starting locations were then determined, and the next trial was run. In general, we ran 100 trials for each value of the independent variable. The performance of our system was measured as the average capture times for each evader in the graph. (Capture times are measured in timesteps; every agent is allowed to move and replan during each timestep.)

V. SIMULATION RESULTS

Our results confirm our hypothesis that modeling adversarial or semi-adversarial agents as randomly dispersing can result




Fig. 9. Pursuit of a marginally adversarial evader (minimax depth one) on the CDG graph. Experiments like this one can be used to determine S experimentally. In this case, the minimum at $S \approx 4$ indicates the optimal S for minimax adversaries of depth one.

in very high capture times. In addition, we demonstrate the ability to efficiently capture highly skilled evaders with fewer pursuers than the capture number of the graph. Particularly, Figure 7 demonstrates the performance of our algorithm when two pursuers are matched against five strongly adversarial agents. Although the capture number of the museum is three, when $S > 1$, our algorithm exploits the evader’s mistakes and coordinates to capture evaders who have moved into the wings of the graph where the capture number is locally only two, or even in some cases one. The extent to which pursuers are willing to abandon nearby, risky pursuits to help guarantee the capture of another more constrained evader depends on the modeled S . We see that as the pursuers’ correctly model the evader adversarialness, capture times decrease. When $S = 1$, as in [15], the pursuers assume random motion so corraling should not be necessary (according to the model), but in reality, the evaders can escape for long periods of time. When $S = 16$, the pursuers know that they are matched against an adversarial evader, and change their behavior accordingly.

We additionally validate our adversarial model on a 100 sample evader in Figure 8. This trial shows similar trends to Figure 7, although minimax is able to avoid capture for longer. The decrease in capture time as the evaders are modeled more correctly clearly shows the reduction in capture time that our model enables over previous methods which assumed $S = 1$.

As discussed in Section I, incorrect assumptions about the adversarialness of an evader can lead to suboptimal capture times. In this case, if S is under- or over-estimated, we expect capture times to increase. We demonstrate how to discover S for a given adversary in Figure 9. In this trial, seven mildly adversarial agents (minimax depth one) are pitted against pursuers who assume varying S . As we expect, capture times reach a minimum when S is correctly estimated (about 4.5).

We also successfully ran trials with 5 evaders and 5 pursuers, showing that our algorithm successfully scales to many agents. This limit was not constrained by our method, but by the fact that on the graphs we present, where the search number is only 3, it is not particularly useful to saturate the

number of pursuers. On larger, more complex graphs, we anticipate being able to scale to a larger number of agents. This contrasts with prior adversarial work, which would have only been able to handle one or two evaders with five pursuers [3].

Finally, for comparison to previous work, we show our adversarial model applied to the search case in Figure 10. In the search variant of the pursuit problem, the agents are pursuing an evader who initially has an equal chance of being anywhere on the graph. As the game evolves, some areas in the graph will have higher probabilities of containing an evader than others. When the random evader model ($S = 1$) is used, pursuers will look first in areas of high probability. Furthermore, the discount factor influences how directly the pursuers will move towards these areas. Since capture only requires one pursuer, the searchers will split up to each visit separate areas of higher probability. As S grows larger, the searchers are more likely to ignore the penalties imposed by the discount factor on future rewards when such plans offer a higher chance of success. That is, a searcher will be more likely to take a longer route to areas of higher probability if it allows coordination with another searcher to jointly increase the probability of capture along the way. The manifestation of this trade-off is shown in Figure 10, where a decrease in capture time is shown by agents who explicitly model the evaders as adversarial, and as a result take a more conservative approach to searching the graph. The improvement of modeling evaders as adversarial is much less pronounced in the search variant since accurate estimates of the evaders location are not available at each timestep. In the pursuit variant, these updates help compensate the random motion model to correctly model adversarial agents.




Fig. 10. Four pursuers search for four evaders whose location is unknown. Evaders are only mildly adversarial in this test, sampling only two trajectories as opposed to the 100 samples drawn by evaders exhibited in our other experiments.

VI. ROBOT RESULTS

In addition to our simulated experiments, we also implemented our algorithm on a three wheeled indoor/outdoor robot platform (see Figure 11a). We emulated a MAGIC 2010-style situation where robots are required to “capture” humans executing fixed patrols. We conducted the experiment in a

Trial	Evader	Captured	Captured by
1	E_a	Step 7	P_a
1	E_b	Step 5	P_b
2	E_a	Step 3	P_a
2	E_b	Step 4	P_b
3	E_a	Step 7	P_b
3	E_b	Step 8	P_a

TABLE I

CAPTURE TIMES FOR PURSUIT-EVADE IMPLEMENTED ON ROBOT PLATFORMS. THE ROBOTS (P_a AND P_b) STARTED IN NODE 1 AND 6 RESPECTIVELY, WHILE THE HUMANS (E_a AND E_b), STARTED IN $\{2, 11, 8\}$ AND $\{7, 12, 3\}$, RESPECTIVELY.

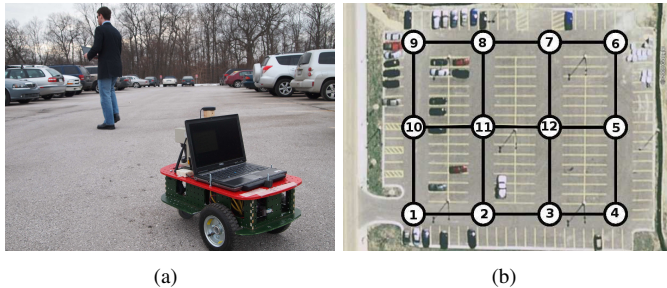


Fig. 11. A robot pursuing a human equipped with a GPS tracking device in a parking lot scenario.

80 by 70 meter parking lot. The area was discretized into a 3x4 square grid of nodes (see Figure 11b). Each of the two pursuers (robots) modeled the two human evaders as $S = 4$. The humans executed fixed patrols on the left two and right two columns, respectively. The robots were equipped with 900MHz radios for communication, a SICK LIDAR for obstacle avoidance, and GPS, IMU and wheel encoders for position estimation.

In total we conducted three trials with varying starting positions for the humans. To ensure that the humans did not have a speed advantage over the robots, an agent was only allowed to move one segment per time step. The results confirm that we are able to efficiently force capture in a real world environment (see Table I).

VII. CONCLUSION AND FUTURE WORK

We introduced a novel method of coping with adversarial agents without necessitating a cost prohibitive adversarial search. Particularly, we demonstrated an efficient, robust and probabilistic pursuit strategy for use against multiple adversarial evaders that does not require partitioning of the pursuers such as in [3, 4].

Our approach scales well as the number of total agents increases, irrespective of the ratio of evaders to pursuers. Although the method does not guarantee optimality, our results show promise in catching evaders who exhibit varying adversarial characteristics, which previous methods were unable to handle. Furthermore, our method is simple in that it has only a very small number of tunable parameters, each of which has a clear meaning.

Although the probabilistic pursuit evasion problem has been studied in depth, there are still open problems in this domain.

Currently, the cunning of the evaders must be known a priori by the strategy and is assumed to be static. For example, changing S dynamically, also on a per-evader basis, would allow better adaptability when the adversarialness of evaders is neither static nor homogeneous. In this work we considered the pursuit variant, where the location of the adversaries is known exactly. Prior work on pursuit with noisy or inaccurate evader position updates could also be expanded to the case of adversarial evaders. Additionally, we have not considered the case where a team of adversarial evaders are “colluding” to escape.

REFERENCES

- [1] “Magic: Multi autonomous ground-robotic international challenge,” <http://www.dsto.defence.gov.au/MAGIC2010/>.
- [2] S. LaValle, D. Lin, L. Guibas, J.-C. Latombe, and R. Motwani, “Finding an unpredictable target in a workspace with obstacles,” in *IEEE International Conference on Robotics and Automation*, 1997.
- [3] M. Vieira, R. Govindan, and G. S. Sukhatme, “Scalable and practical pursuit-evasion with networked robots,” *Journal of Intelligent Service Robotics. Special Issue on Networked Robots*, 2009.
- [4] G. Hollinger, S. Singh, and A. Kehagias, “Efficient, guaranteed search with multi-agent teams,” in *Robotics: Science and Systems*, 2009.
- [5] G. Hollinger and S. Singh, “Proofs and experiments in scalable, near-optimal search by multiple robots,” in *Robotics: Science and Systems*, 2008.
- [6] J. O’Rourke, *Art Gallery Theorems and Algorithms*, 1987.
- [7] L. Shenato, S. Oh, S. Sastry, and P. Bose, “Swarm coordination for pursuit evasion games using sensor networks,” in *ICRA*, 2005.
- [8] T. Parsons, *Pursuit-Evasion in a Graph*, 1976, pp. 426–441.
- [9] P. Seymour and R. Thomas, “Graph searching, and a min-max theorem for tree-width,” *J. Combin. Theory Ser. B*, pp. 22–33.
- [10] V. Isler, S. Kannan, and S. Khanna, “Randomized pursuit-evasion in a polygonal environment,” in *IEEE Transactions on robotics*, 2005.
- [11] L. E. Parker, “Distributed algorithms for multi-robot observation of multiple moving targets,” *Autonomous Robotics*, pp. 1–30, 2002.
- [12] J. P. Hespanha, H. J. Kim, and S. Sastry, “Multiple-agent probabilistic pursuit-evasion games,” in *38th Conf. on Decision and Control*, 1999.
- [13] R. Vidal, O. Shakernia, H. J. Kim, H. Shim, and S. Sastry, “Probabilistic pursuit evasion games: Theory, implementation and experimental evaluation,” in *IEEE Transactions on Robotics and Automation*, 2002, pp. 100–107.
- [14] A. Kehagias, G. Hollinger, and S. Singh, “A graph search algorithm for indoor pursuit/evasion,” vol. 50, pp. 1305–1317, 2009.
- [15] G. Hollinger, J. Djughash, and S. Singh, “Coordinated search in cluttered environments using range from multiple robots,” *Field and Service Robotics*, 2008.
- [16] G. Hollinger, A. Kehagias, and S. Singh, “Probabilistic strategies for pursuit in cluttered environments with multiple robots,” in *IEEE International Conference on Robotics and Automation*, April 2007.
- [17] M. Kearns, Y. Mansour, and A. Y. Ng, “A sparse sampling algorithm for near-optimal planning in large markov decision processes,” *Machine Learning*, vol. 49, pp. 193–208, 2002.
- [18] R. Nair, M. Tambe, M. Yokoo, D. Pynadath, and S. Marsella, “Taming decentralized pomdps: Towards efficient policy computation for multiagent settings,” in *International Joint Conference On Artificial Intelligence*, 2003, pp. 705–711.
- [19] J. Zhen, H. Yu, W. Liang, and P. Zeng, “Randomized strategies to coordinate multiple robotic pursuers in the pursuit of an adversarial evaders,” in *1th International Federation of Automatic Control*, 2008.