Supplementary Material: Path Planning for Information Gathering with Lethal Hazards and No Communication

Michael Otte[1] and Donald Sofge[2]

[1] University of Maryland, College Park, MD 20742
[2] U.S. Naval Research Laboratory, Washington DC 20375

Abstract. This document contains additional supplementary material in support of our manuscript which has been submitted to the Workshop on the Algorithmic Foundations of Robotics. Included are: additional figures for the experiment presented in the main document, additional experimental results, and captions for videos. The videos themselves appear as additional supplementary material.
In the main document we consider a scenario where agents search for targets in a hazardous environment that prevents communication. Agents in the field cannot communicate, and hazards are only directly observable by the agents that are destroyed by them. Thus, beliefs about hazard locations must be inferred by sending agents to travel along various paths and then observing which agents survive. In other words, agent survival along a path can be used as a sensor for hazard detection; we call this form of sensor a "path-based sensor". We present a recursive Bayesian update for path-based sensors, and leverage it to calculate the expected information gained about both hazards and targets along a particular path. We formalize the resulting iterative information based path planning problem that results from this scenario, and present an algorithm to solve it. Agents iteratively foray into the field. The next path each agent follows is calculated to maximize a weighted combination of the expected information gained about targets and hazards (where the weighting is defined by user preferences). The method is evaluated in Monte Carlo simulations, and we observe that it outperforms other techniques.

1 Introduction

Section 2 contains information regarding the video that is uploaded as part of the supplementary material. Section 3 contains additional figures for the experiment presented in the main document. Section 4 contains figures for another experiment that has less targets and hazards.

2 Video

A video appearing as part of the supplementary materials show random trials using same environment as the experiment presented in our main paper, but for adversary lethality rates of 0.5.

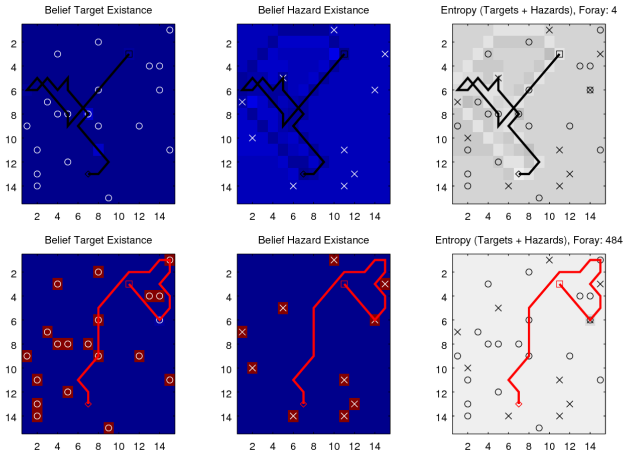

Fig. 1: Example of what appears in the video that is included in supplementary material. The belief of hazard existence and target existence in the left and center sub-plots, respectively. Likelihood is drawn on a color scale where blue is low likelihood and red is high likelihood. The locations of targets and hazards is initially unknown to the system. In the plots, the locations with hazards are marked with an 'x' and the locations of targets are marked with an 'o'. The entropy (or summed entropy) associated with the objective function are appear in the left sub-plot, and drawn in gray-scale where darker represents more entropy. The path that maximizes the weighted objective over target and entropy is depicted in black in all plots (Top). Video frames alternate between planned paths, and the update that occurs after a a path has been attempted. Paths turn red during updates for iterations in which agents are destroyed (Bottom).

The video shows a sequence of clips, where each clip is for a different parameter selection or comparison method. The belief in adversaries and targets in the left and center sub-plots, respectively. The entropy (or summed entropy) associated with the planning objective are appear in the left sub-plot. Video frames alternate between planned paths, and the update that occurs after a a path has been attempted. Paths turn red for iterations in which agents are destroyed. An example appears in Figure 1.

3 Additional Figures For the Experiment Presented in the Main Document

In this section we include additional experimental results for the experiment presented in the main document. The first three subsections contain results related to the method we present. The final three subsections show results for the comparison methods. Each method or parameter combination is tested for various adversary lethality rates.

Much of the text in this section has been taken from the main document.

We use Monte Carlo trials in simulation to evaluate the performance of our method, and compare it to other approaches. The environment is represented by a 15×15 grid map. Movement is defined by a 9-grid of connectivity (8-grid neighbors plus self transitions) where each move takes the agent one time-step further in time. The agent has enough endurance to make 25 moves. Agent malfunction rate is $p_{\text{malfunc}} = .01$ per time step (thus, on average agents arbitrarily malfunction in 1/4 of all forays of length 25). Additional experiments are included in the supplementary materials. In each trial the start and goal uplink points are placed uniformly at random. 10 non-start/goal locations are picked uniformly at random (no replacement) and populated with hazards. This is repeated for 20 non-start/goal locations that are populated with targets.

We test our algorithm using three different objectives: weighting information from targets and hazards equally $c_X, c_Z = \{1, 1\}$; gathering only target information $c_X, c_Z = \{1, 0\}$; and gathering only hazard information $c_X, c_Z = \{0, 1\}$. We compare to three other ideas: (1) 1-step look ahead information surfing; (2) a Markov random walk; and (3) planning paths to gather target information while ignoring hazards altogether (by not accounting for the probability of being destroyed when evaluating the expected information gain, and assuming a $c_X, c_Z = \{1, 0\}$ objective).

Our method and information surfing both track and update target and obstacle beliefs, and use the probability of hazard existence to weight the expected information that will be gained about targets and/or hazards. The path of the random walk is calculated before the agent departs such that the resulting-path sensor can be used to infer hazard presence based on whether or not it survives. In all methods agent movement is only allowed in directions from which the agent can still reach the goal given its fuel reserves.

Each of the following subsections presents results for a particular method. Each method is tested across a variety of hazard lethality ratios including: .01, .2, .4, .6, .8, and .99.

3.1 Info Path, $C = [1, 1]$, 20 Targets and 10 Hazards

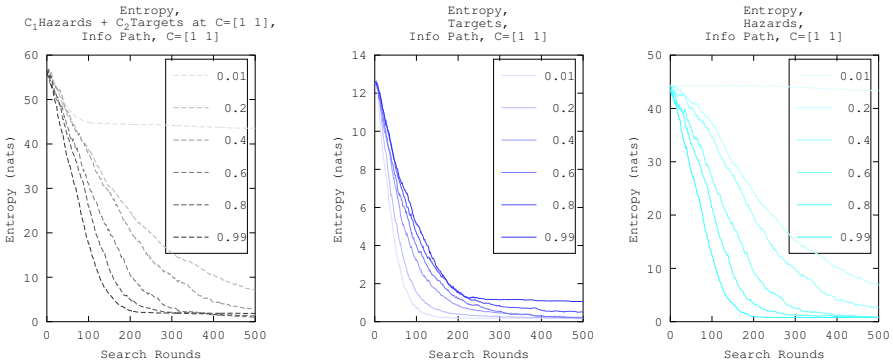

Fig. 2: Hazard and Target entropies vs. search round (mean over 30 random trials) when hazards have a 60% kill ratio. Left: an equally weighted ($C = [1, 1]$) combination of hazard and target entropy. Center: Target entropy. Right: Hazard entropy.

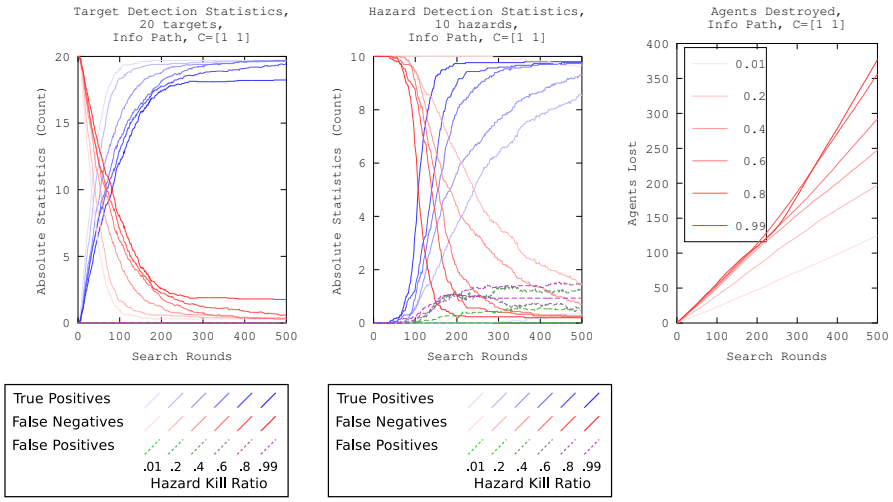

Fig. 3: Left: Target statistics. Center: Hazard statistics. Right: total agent losses over time. Statistics include true/false positives and false negatives of target locations. Likelihood $\geq .95$ is defined as a positive detection. All plots show mean results over 30 random trials. In these experiments hazards have a 60% kill success rate. (Agents are expendable such that any costs associated with their losses are negligible compared to the information that is gained from their loss).

3.2 Ignore Adversary Info (but use survival probability), 20 Targets and 10 Hazards

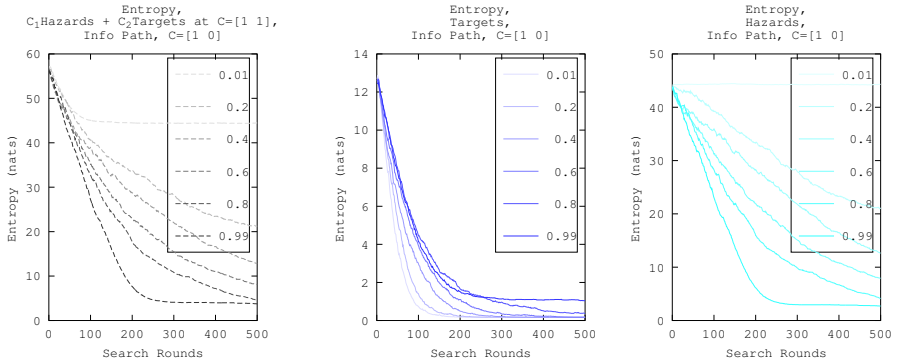

Fig. 4: Hazard and Target entropies vs. search round (mean over 30 random trials) when hazards have a 60% kill ratio. Left: an equally weighted ($C = [1, 1]$) combination of hazard and target entropy. Center: Target entropy. Right: Hazard entropy.

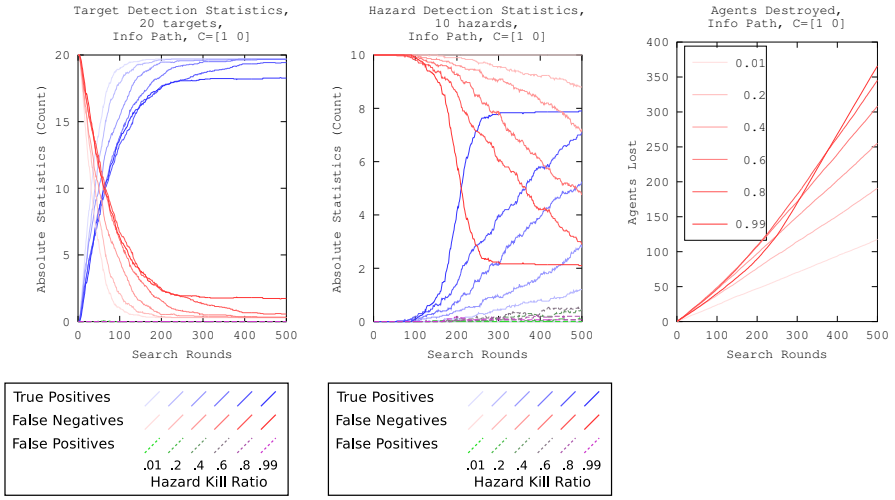

Fig. 5: Left: Target statistics. Center: Hazard statistics. Right: total agent losses over time. Statistics include true/false positives and false negatives of target locations. Likelihood $\geq .95$ is defined as a positive detection. All plots show mean results over 30 random trials. In these experiments hazards have a 60% kill success rate. (Agents are expendable such that any costs associated with their losses are negligible compared to the information that is gained from their loss).

3.3 Ignore Target Info, 20 Targets and 10 Hazards

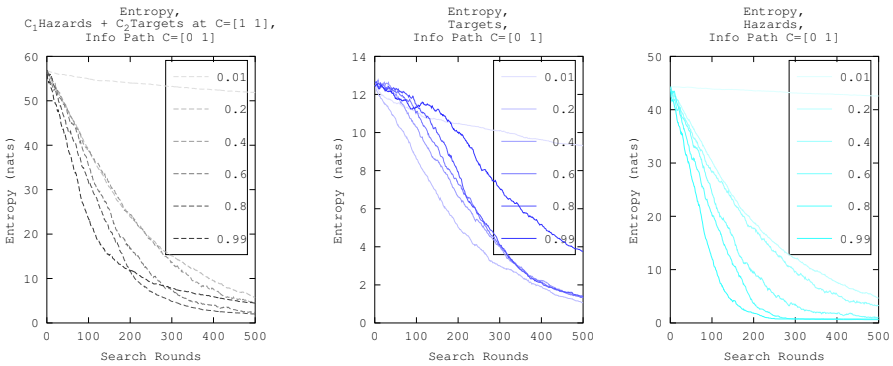

Fig. 6: Hazard and Target entropies vs. search round (mean over 30 random trials) when hazards have a 60% kill ratio. Left: an equally weighted ($C = [1, 1]$) combination of hazard and target entropy. Center: Target entropy. Right: Hazard entropy.

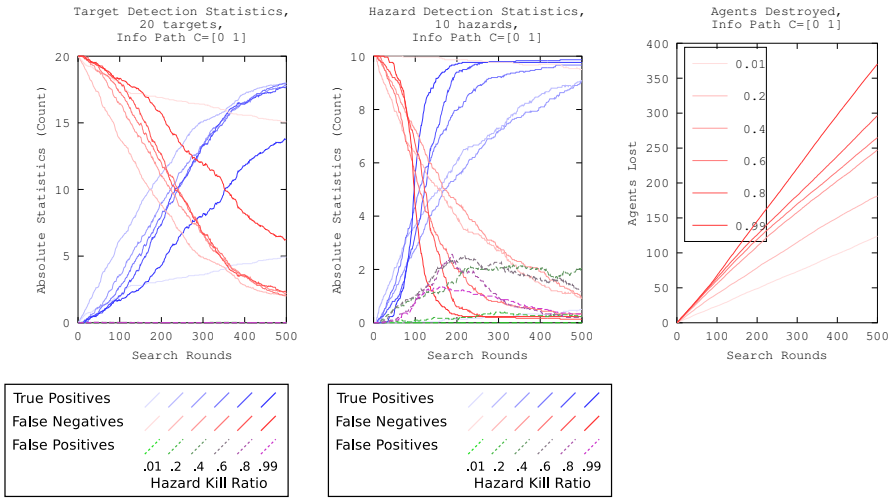

Fig. 7: Left: Target statistics. Center: Hazard statistics. Right: total agent losses over time. Statistics include true/false positives and false negatives of target locations. Likelihood $\geq .95$ is defined as a positive detection. All plots show mean results over 30 random trials. In these experiments hazards have a 60% kill success rate. (Agents are expendable such that any costs associated with their losses are negligible compared to the information that is gained from their loss).

3.4 Information Surfing (1-Step), 20 Targets and 10 Hazards

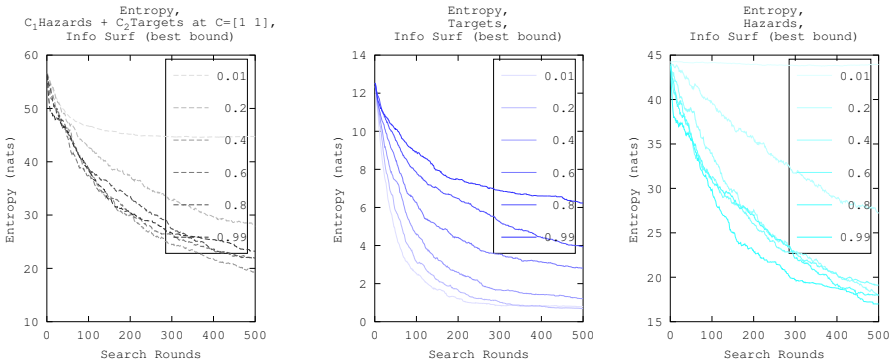

Fig. 8: Hazard and Target entropies vs. search round (mean over 30 random trials) when hazards have a 60% kill ratio. Left: an equally weighted ($C = [1, 1]$) combination of hazard and target entropy. Center: Target entropy. Right: Hazard entropy.

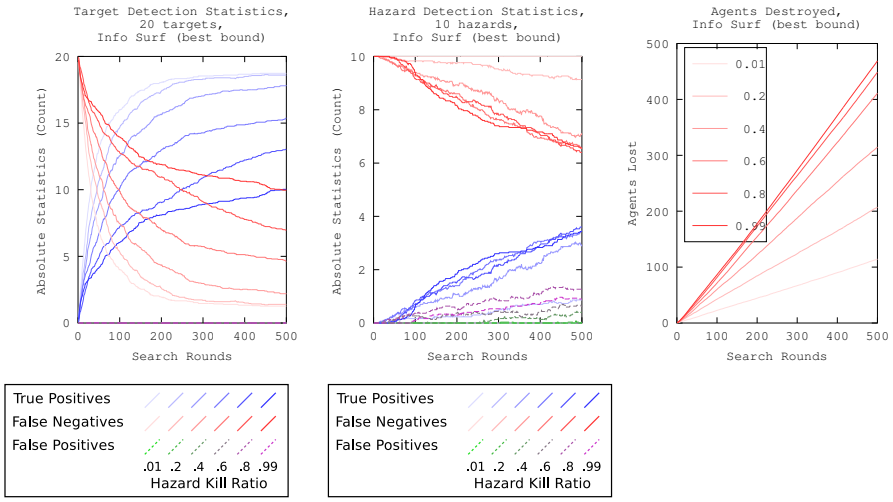

Fig. 9: Left: Target statistics. Center: Hazard statistics. Right: total agent losses over time. Statistics include true/false positives and false negatives of target locations. Likelihood $\geq .95$ is defined as a positive detection. All plots show mean results over 30 random trials. In these experiments hazards have a 60% kill success rate. (Agents are expendable such that any costs associated with their losses are negligible compared to the information that is gained from their loss).

3.5 Markov Random Walk, 20 Targets and 10 Hazards

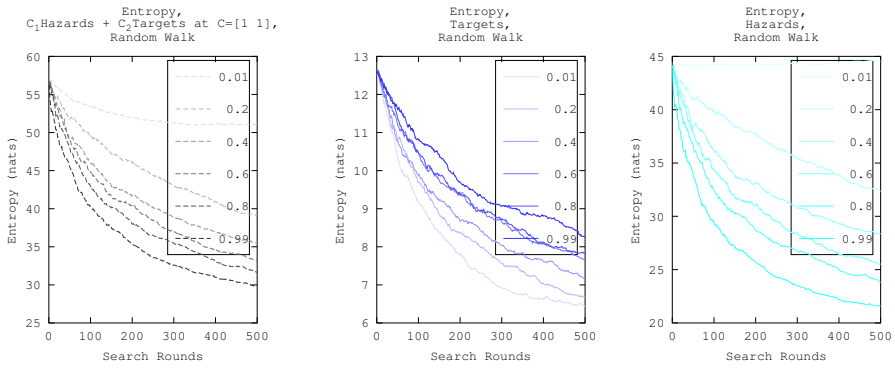

Fig. 10: Hazard and Target entropies vs. search round (mean over 30 random trials) when hazards have a 60% kill ratio. Left: an equally weighted ($C = [1, 1]$) combination of hazard and target entropy. Center: Target entropy. Right: Hazard entropy.

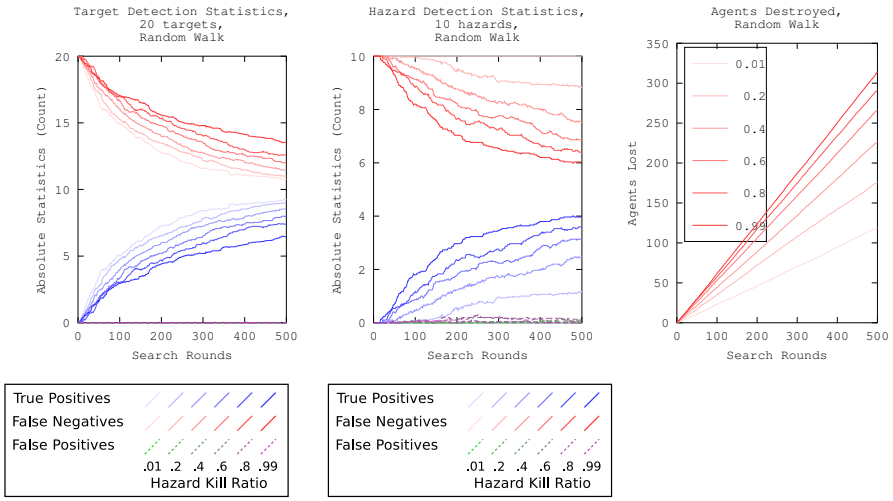

Fig. 11: Left: Target statistics. Center: Hazard statistics. Right: total agent losses over time. Statistics include true/false positives and false negatives of target locations. Likelihood $\geq .95$ is defined as a positive detection. All plots show mean results over 30 random trials. In these experiments hazards have a 60% kill success rate. (Agents are expendable such that any costs associated with their losses are negligible compared to the information that is gained from their loss).

3.6 Ignore Adversaries Completely, 20 Targets and 10 Hazards

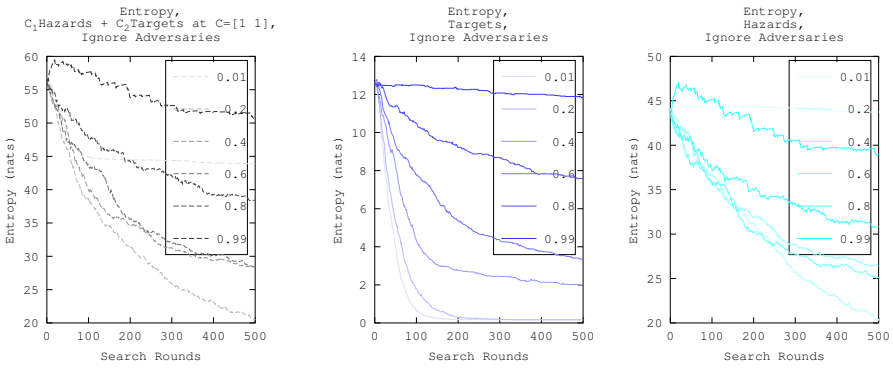

Fig. 12: Hazard and Target entropies vs. search round (mean over 30 random trials) when hazards have a 60% kill ratio. Left: an equally weighted ($C = [1, 1]$) combination of hazard and target entropy. Center: Target entropy. Right: Hazard entropy.

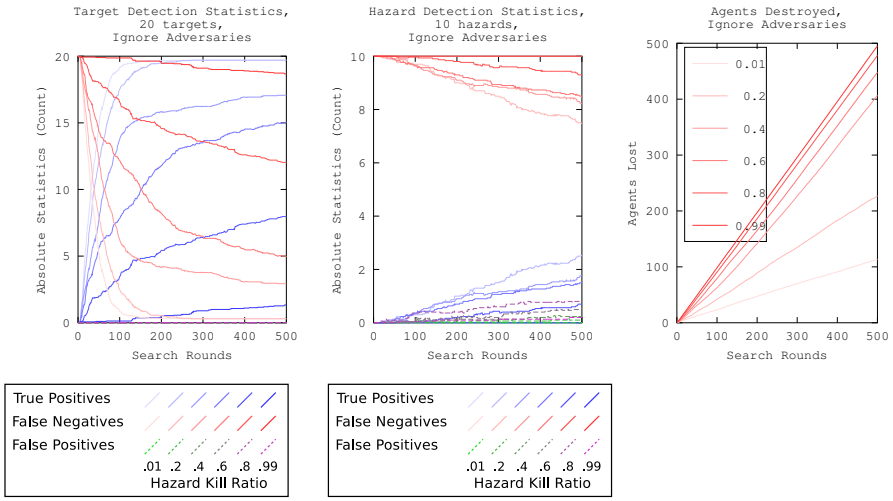

Fig. 13: Left: Target statistics. Center: Hazard statistics. Right: total agent losses over time. Statistics include true/false positives and false negatives of target locations. Likelihood $\geq .95$ is defined as a positive detection. All plots show mean results over 30 random trials. In these experiments hazards have a 60% kill success rate. (Agents are expendable such that any costs associated with their losses are negligible compared to the information that is gained from their loss).

4 Experiment with 4 targets and 6 hazards

In this section we include an additional experiment with 4 targets and 6 hazards. The first three subsections contain results related to the method we present. The final three subsections show results for the comparison methods. Each method or parameter combination is tested for various adversary lethality rates.

The format of these experiments and trials is identical to those presented in the previous section, except that start and goal locations are the same across all random trials (and also that the number of hazards and targets are 4 and 6 respectively).

4.1 Info Path, $C = [1, 1]$, 4 Targets and 6 Hazards

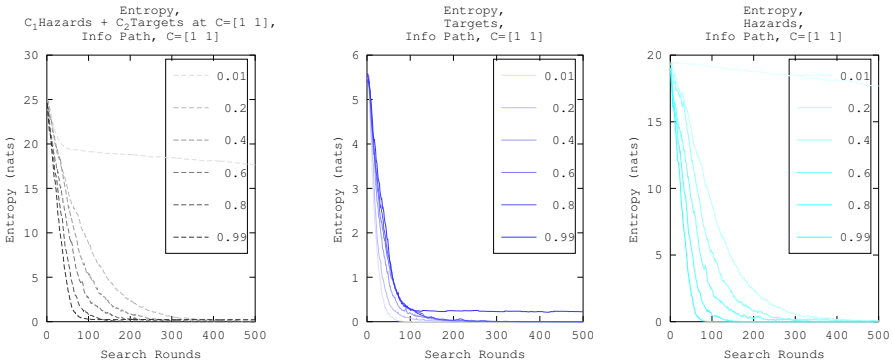

Fig. 14: Hazard and Target entropies vs. search round (mean over 30 random trials) when hazards have a 60% kill ratio. Left: an equally weighted ($C = [1, 1]$) combination of hazard and target entropy. Center: Target entropy. Right: Hazard entropy.

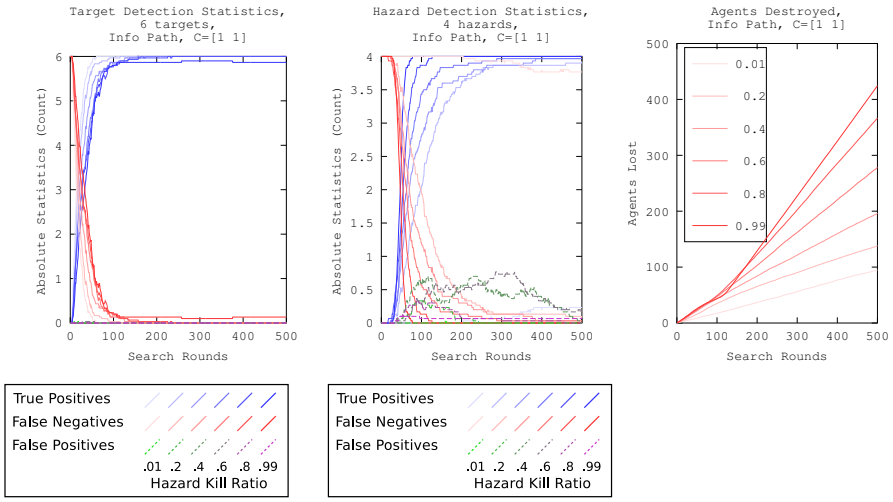

Fig. 15: Left: Target statistics. Center: Hazard statistics. Right: total agent losses over time. Statistics include true/false positives and false negatives of target locations. Likelihood $\geq .95$ is defined as a positive detection. All plots show mean results over 30 random trials. In these experiments hazards have a 60% kill success rate. (Agents are expendable such that any costs associated with their losses are negligible compared to the information that is gained from their loss).

4.2 Ignore Adversary Info (but use survival probability), 4 Targets and 6 Hazards

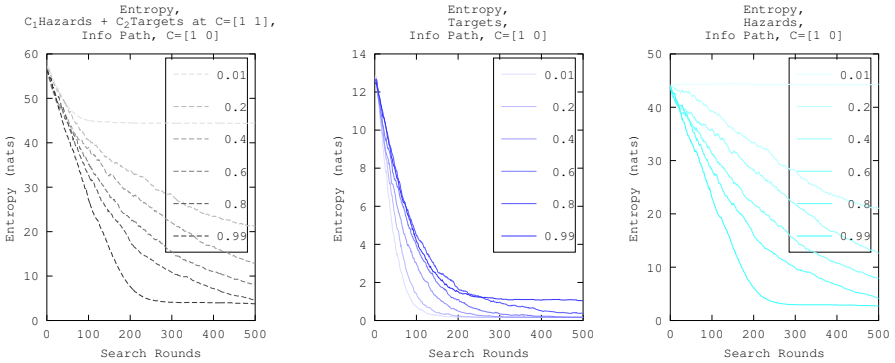

Fig. 16: Hazard and Target entropies vs. search round (mean over 30 random trials) when hazards have a 60% kill ratio. Left: an equally weighted ($C = [1, 1]$) combination of hazard and target entropy. Center: Target entropy. Right: Hazard entropy.

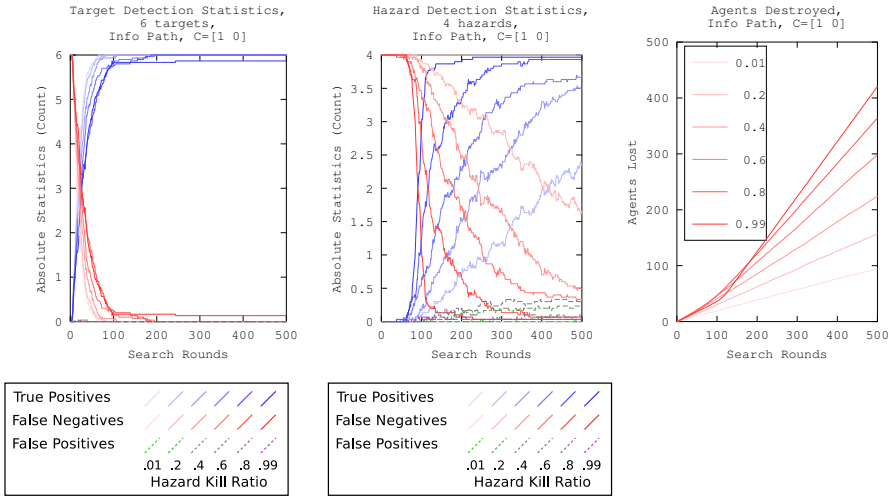

Fig. 17: Left: Target statistics. Center: Hazard statistics. Right: total agent losses over time. Statistics include true/false positives and false negatives of target locations. Likelihood $\geq .95$ is defined as a positive detection. All plots show mean results over 30 random trials. In these experiments hazards have a 60% kill success rate. (Agents are expendable such that any costs associated with their losses are negligible compared to the information that is gained from their loss).

4.3 Ignore Target Info, 4 Targets and 6 Hazards

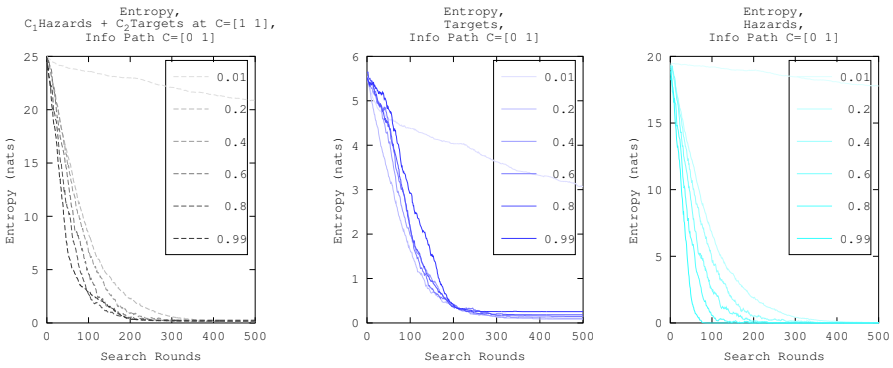

Fig. 18: Hazard and Target entropies vs. search round (mean over 30 random trials) when hazards have a 60% kill ratio. Left: an equally weighted ($C = [1, 1]$) combination of hazard and target entropy. Center: Target entropy. Right: Hazard entropy.

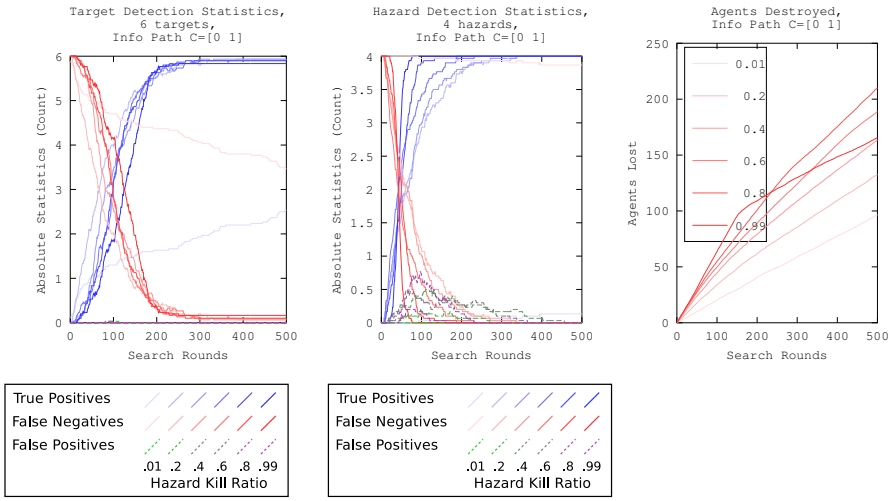

Fig. 19: Left: Target statistics. Center: Hazard statistics. Right: total agent losses over time. Statistics include true/false positives and false negatives of target locations. Likelihood $\geq .95$ is defined as a positive detection. All plots show mean results over 30 random trials. In these experiments hazards have a 60% kill success rate. (Agents are expendable such that any costs associated with their losses are negligible compared to the information that is gained from their loss).

4.4 Information Surfing (1-Step), 4 Targets and 6 Hazards

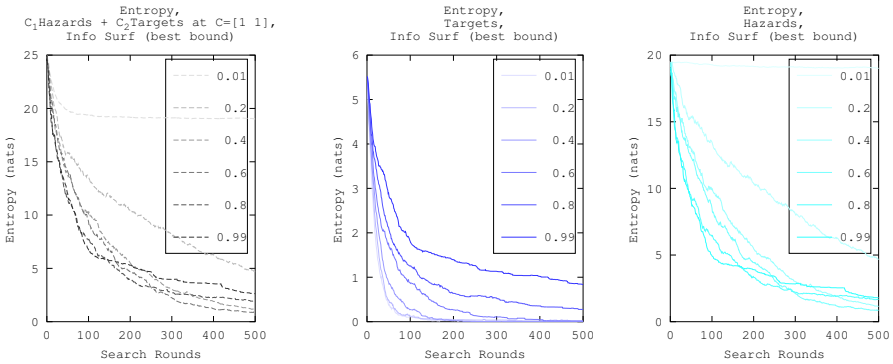

Fig. 20: Hazard and Target entropies vs. search round (mean over 30 random trials) when hazards have a 60% kill ratio. Left: an equally weighted ($C = [1, 1]$) combination of hazard and target entropy. Center: Target entropy. Right: Hazard entropy.

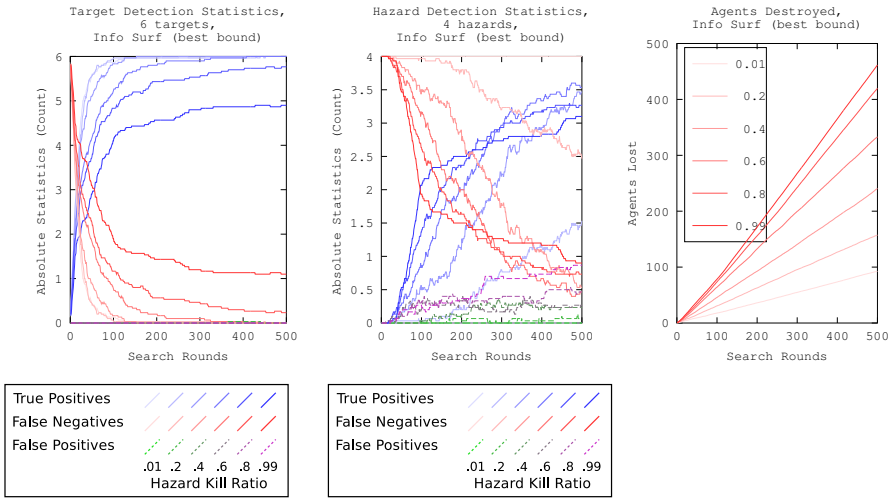

Fig. 21: Left: Target statistics. Center: Hazard statistics. Right: total agent losses over time. Statistics include true/false positives and false negatives of target locations. Likelihood $\geq .95$ is defined as a positive detection. All plots show mean results over 30 random trials. In these experiments hazards have a 60% kill success rate. (Agents are expendable such that any costs associated with their losses are negligible compared to the information that is gained from their loss).

4.5 Markov Random Walk, 4 Targets and 6 Hazards

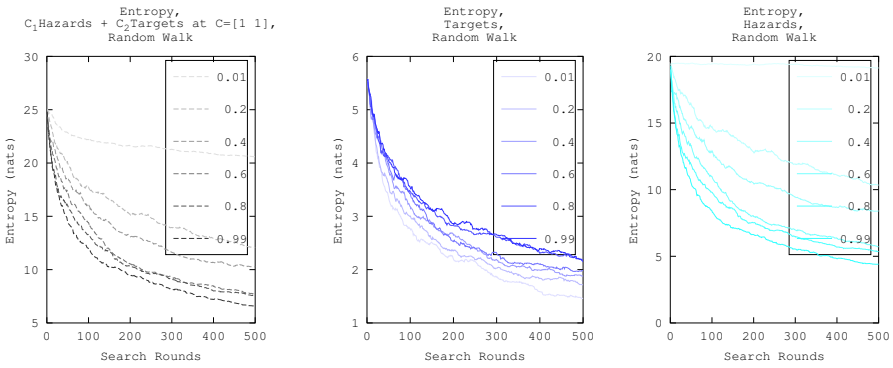

Fig. 22: Hazard and Target entropies vs. search round (mean over 30 random trials) when hazards have a 60% kill ratio. Left: an equally weighted ($C = [1, 1]$) combination of hazard and target entropy. Center: Target entropy. Right: Hazard entropy.

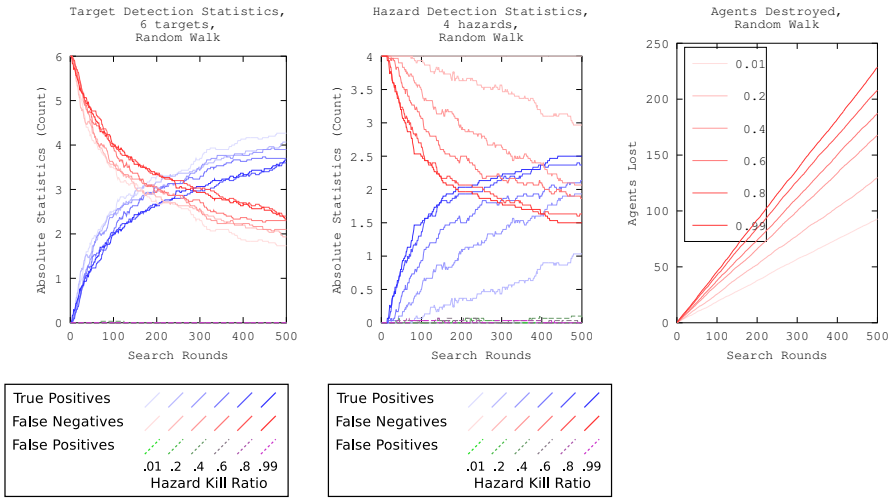

Fig. 23: Left: Target statistics. Center: Hazard statistics. Right: total agent losses over time. Statistics include true/false positives and false negatives of target locations. Likelihood $\geq .95$ is defined as a positive detection. All plots show mean results over 30 random trials. In these experiments hazards have a 60% kill success rate. (Agents are expendable such that any costs associated with their losses are negligible compared to the information that is gained from their loss).

4.6 Ignore Adversaries Completely, 4 Targets and 6 Hazards

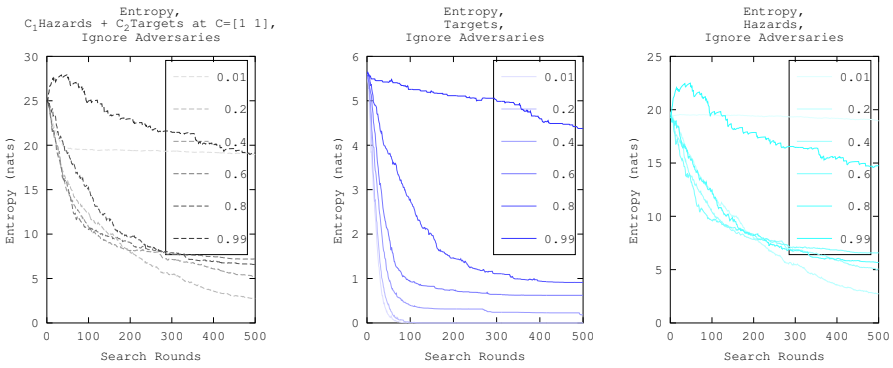

Fig. 24: Hazard and Target entropies vs. search round (mean over 30 random trials) when hazards have a 60% kill ratio. Left: an equally weighted ($C = [1, 1]$) combination of hazard and target entropy. Center: Target entropy. Right: Hazard entropy.

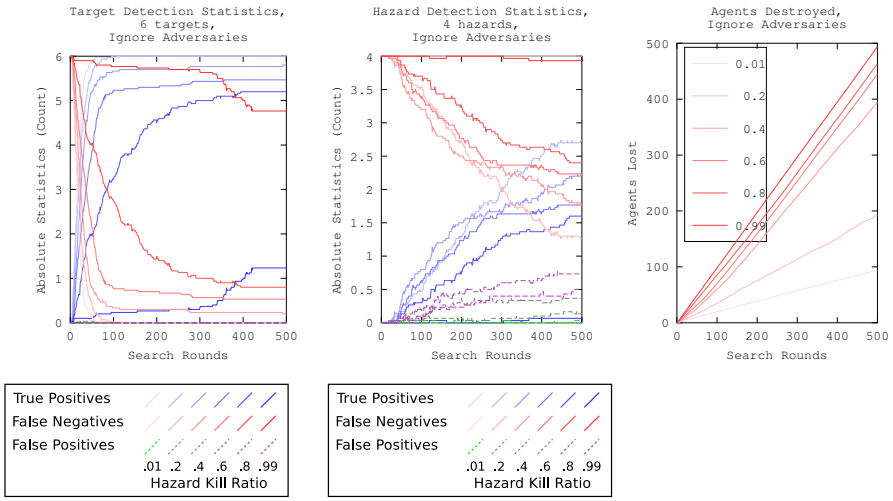

Fig. 25: Left: Target statistics. Center: Hazard statistics. Right: total agent losses over time. Statistics include true/false positives and false negatives of target locations. Likelihood $\geq .95$ is defined as a positive detection. All plots show mean results over 30 random trials. In these experiments hazards have a 60% kill success rate. (Agents are expendable such that any costs associated with their losses are negligible compared to the information that is gained from their loss).