Robotics and Autonomous Systems 86 (2016) 1–12

Contents lists available at ScienceDirect

Robotics and Autonomous Systems

journal homepage: www.elsevier.com/locate/robot

Multi-agent motion planning using Bayes risk[☆]

Matthew J. Bays [a,*], Apoorva Shende [b], Daniel J. Stilwell [b,**]

[a] Naval Surface Warfare Center, Panama City Division, Panama City, FL 32407, United States

[b] Bradley Department of Electrical & Computer Engineering, Virginia Tech, Blacksburg, VA 24060, United States

HIGHLIGHTS

- We introduce an approach to multi-agent motion control.
- Approach jointly minimizes a cost function utilizing Bayes risk for classification.
- We apply the framework in the context of target interception and collision avoidance.
- Framework uses a particle approach combined with mixed integer linear programming.

ARTICLE INFO

Article history:
Received 8 December 2014
Received in revised form
18 April 2016
Accepted 22 July 2016
Available online 29 August 2016

Keywords:
Autonomous agents
Robotics
Motion planning
Bayes risk
Classification
Decision theory

ABSTRACT

We introduce a novel approach to controlling the motion of a team of agents so that they jointly minimize a cost function utilizing Bayes risk. Bayes risk is a useful measure of performance for applications where agents must perform a classification task, but is often difficult to compute analytically for many applications involving agent state variables. We use a particle-based approach that allows us to approximate Bayes risk and express the optimization problem as a mixed-integer linear program. By minimizing Bayes risk, agents are able to account explicitly for the costs associated with correct and incorrect classification. We illustrate our approach with a target interception problem in which a team of mobile agents must intercept mobile targets that are likely to enter a specified area in the near future. We show that the cooperative agent motion that minimizes a cost function utilizing Bayes risk is an efficient way to achieve selective interception.

1. Introduction

We investigate a method for coordinating the motion of a group of agents or mobile sensors to improve performance of a classification task. We propose a particle-based approximation of Bayes risk in order to develop a cost function to guide the motion of the agents. Bayes risk is the prior-averaged cost incurred by a decision rule [1]. Bayesian probabilistic techniques are particularly relevant for classification-based feedback control since they inherently update the probabilistic distribution of the underlying state as new observations become available. The updated state distribution that is computed using Bayesian filtering techniques can be used to compute the Bayes risk corresponding to the current observations. Prior work on cooperative control has often addressed optimal sensor motion for improved sensing performance, as in the case of maximizing information [2]. In typical cases, sensor motions produce optimal state estimates, but would not necessarily lead to optimal, or even improved, classification performance. The research into the decision-making and classification has largely been the work of the statistics and signal processing community, where an optimal decision rule is generated for the observations already gathered. A good example of these contributions in the area of target classification is [3]. The primary contribution of our work is attempting to bridge the gap between classification and agent motion control (e.g. cooperative control as in [4]) by proposing a methodology for which cooperative sensor motion is used to intentionally improve classification performance. We provide context for our Bayes risk-based motion control using a target interception application.

Bayes risk has been used frequently in the signal processing community; see for example [5–8]. However, its use in motion

☆ This paper was presented in part at the *IEEE International Conference on Robotics and Automation*, St. Paul, MN, May 2012.

* Corresponding author. Fax: +1 561 818 1765.
** Corresponding author. Fax: +1 540 231 3362.
E-mail addresses: mjb222@ieee.org (M.J. Bays), apoorva@vt.edu (A. Shende), stilwell@vt.edu (D.J. Stilwell).

http://dx.doi.org/10.1016/j.robot.2016.07.009
0921-8890/

planning has been limited. Blackmore et al. in [9] develop an approach for model selection from a family of dynamic systems using a special case of Bayes risk that is equivalent to the probability of error. In their approach, control signals for the dynamic system are computed that minimize an upper bound for the probability of error. We generate control signals that affect the paths for a team of mobile agents. However, we choose the control signals to minimize an approximation to a more general form of Bayes risk.

The problem of area protection has been studied using a variety of techniques. Smith uses formal geometric approaches in [10] to choose paths for vehicles servicing randomly generated targets whose locations are deterministically known. In contrast, our particle-based solution to target interception takes into account uncertainty in target locations at every time-step. Agmon develops a novel patrol scheme in order to minimize the probability of adversary penetration in [11]. Our work complements Agmon's by addressing the same scenario, but optimizing the agents' reactions to incoming targets instead of optimizing the standard patrol pattern. Beard uses a k-best paths graph search to intercept deterministic targets within an area in [12]. Like our work, Earl and D'Andrea studied the target interception problem in a MILP framework [13]. However, we introduce a formal risk metric, and generalize system dynamics from a deterministic framework to incorporate stochastic dynamics and noise. Additionally, we formulate the problem using a receding-horizon control approach, limiting the planning horizon to a subset of the total mission time and thus complexity of the problem. Chasparis and Shamma solve the problem of area protection in [14] using a discrete-time, discrete-space resource allocation framework. Control policies are modeled as the transition of resources from cell to cell. The objective of maximizing the interceptions of enemy vehicles is modeled using a linear cost function. In contrast, our work explicitly addresses Bayesian risk. Shende solves the problem of stochastic target interception using closed-form properties of the Gaussian distribution in [15] for instances where the system noise is normally distributed, however does not use Bayes risk as a cost function.

The use of Bayes risk for collision avoidance has been previously demonstrated by Jansson and Gustafsson for automobile applications in [16]. In their work, a binary decision rule was implemented. Their decision rule was optimal in the Bayes risk sense, and if met, a collision avoidance maneuver was executed. Our work explicitly couples agent motion control to a cost function utilizing Bayes risk, and as a result, minimizes the risk of a collision.

This paper is outlined as follows. We present the general problem formulation we use for our agent-based motion control and target interception, its relation to Bayes risk, and our framework's underlying assumptions in Section 2. We then propose a particle-based approximation for calculating the necessary probabilities and thereby Bayes risk in Section 3. We extend the particle framework to incorporate collision avoidance in Section 4. Specific estimation techniques used in order to create the particle framework, and a discussion of the potential impact of assumptions on real implementations are provided in Section 5. Finally, we present simulation results of the framework using the interception-only model as well as the collision avoidance extension in Section 6.

2. Problem formulation & Bayes risk

In our application of target interception, M cooperative agents and N targets maneuver within a convex area $G \subset \mathbb{R}^2$. The agents must prevent the targets from entering a subset $A \subset G$ by selectively intercepting those targets that will enter the area of protection. There is uncertainty associated with initial target and agent positions, as well as their dynamics. Additional uncertainty is also associated with target measurements from sensors onboard the agents. The problem is to develop a control law over the *optimization interval* $[t, t+T]$ to guide the cooperative agents such that they intercept all the targets that would have entered the area of protection if not intercepted, and to avoid taking action against targets that would not have entered the area. Additionally, we assume that the targets are not adversarial. That is, the targets do not perform control actions taking into account the states of the agents.

2.1. The target interception problem

We will now define the decision problem with respect to motion planning and classification. At each time $\tau \in [t, t+T]$ within the optimization interval, we consider a target trajectory over the interval $[\tau, \tau + T']$. We will call this horizon the *threat assessment interval*. Note that the threat assessment horizon T' and planning horizon T are independently specified. The trajectory of target i is classified to be within the set of trajectories considered a *threat*, or the set of trajectories considered *not a threat*. The event that the trajectory of target i over the interval $[\tau, \tau + T']$ is a threat is denoted as $\Omega_i(\tau)$. The event that the target trajectory is within the set of trajectories considered not a threat is denoted as $\sim\Omega_i(\tau)$. Target i is classified as a threat if it enters the area of protection A during this interval. A correct decision that should be reflected through the agent control actions is to *intercept* the target. This decision is denoted $\Gamma_i(\tau)$ for target i at time τ. A trajectory that does not enter an area of protection within the time horizon $[\tau, \tau + T']$, or has been intercepted by an agent before τ, is categorized as *not a threat* ($\sim\Omega_i(\tau)$), and corresponds to a correct decision of *non-interception* ($\sim\Gamma_i(\tau)$). For a target, the decision rule that must hold true for the target interception problem may be written as

$$\Omega_i(\tau) \rightarrow \Gamma_i(\tau') \tag{1}$$

$$\sim\Omega_i(\tau) \rightarrow \sim \Gamma_i(\tau') \tag{2}$$

for some time $\tau' < \tau \in [t, t+T]$. However, (1)–(2) is a decision rule that would only work if the problem were purely deterministic. In order to take into account the probabilistic nature of the target interception problem, we turn to Bayes risk.

2.2. Bayes risk for target interception

In the following sections, we describe an agent control law designed to minimize the cumulative Bayes risk associated with the target interception problem involving the threat and interception events defined in (1)–(2). Let $R_{\Omega_i}(\tau)$ be the cost incurred conditioned on hypothesis Ω_i being true at time τ. In Bayes risk terminology, $R_{\Omega_i}(\tau)$ is often called the *conditional risk* associated with the hypothesis $\Omega_i(\tau)$ being true. Similarly, $R_{\sim\Omega_i}(\tau)$ is the conditional risk associated with the hypothesis $\sim\Omega_i(\tau)$ being true at time τ. The Bayes risk at time τ is then given by

$$R_i(\tau) = R_{\Omega_i}(\tau)P(\Omega_i(\tau)) + R_{\sim\Omega_i}(\tau)P(\sim\Omega_i(\tau)), \tag{3}$$

where $P(\Omega_i(\tau))$ and $P(\sim\Omega_i(\tau))$ are the priors of the two classes at time τ. The conditional risks $R_{\Omega_i}(\tau)$ and $R_{\sim\Omega_i}(\tau)$ are written

$$\begin{aligned} R_{\sim\Omega_i}(\tau) &= C_{0,0}P(\sim\Gamma_i(\tau) \mid \sim \Omega_i(\tau)) + C_{1,0}P(\Gamma_i(\tau) \mid \sim \Omega_i(\tau)) \\ R_{\Omega_i}(\tau) &= C_{0,1}P(\sim\Gamma_i(\tau) \mid \Omega_i(\tau)) + C_{1,1}P(\Gamma_i(\tau) \mid \Omega_i(\tau)), \end{aligned} \tag{4}$$

where $C_{0,0}$, $C_{1,0}$, $C_{0,1}$ and $C_{1,1}$ are user-selected coefficients associated with the cost of correct and incorrect classification. From the standard definition of Bayes risk, $C_{0,0}$ and $C_{1,1}$ are the costs of correct classification, and $C_{1,0}$ and $C_{0,1}$ are the costs of false-alarm and missed detection, respectively [1]. If we assume $\Omega_i(\tau)$ and $\Gamma_i(\tau)$ to