Artificial Intelligence 174 (2010) 1142–1171

Contents lists available at ScienceDirect

Artificial Intelligence

www.elsevier.com/locate/artint

Robust solutions to Stackelberg games: Addressing bounded rationality and limited observations in human cognition

James Pita [a], Manish Jain [a], Milind Tambe [a,*], Fernando Ordóñez [a,b], Sarit Kraus [c,d]

[a] University of Southern California, Los Angeles, CA 90089, United States
[b] University of Chile, Santiago, Chile
[c] Bar-Ilan University, Ramat-Gan 52900, Israel
[d] Institute for Advanced Computer Studies, University of Maryland, College Park, MD 20742, United States

ARTICLE INFO

Article history:
Received 24 November 2009
Received in revised form 4 July 2010
Accepted 10 July 2010
Available online 16 July 2010

Keywords:
Behavioral game theory
Security
Stackelberg
Uncertainty

ABSTRACT

How do we build algorithms for agent interactions with human adversaries? Stackelberg games are natural models for many important applications that involve human interaction, such as oligopolistic markets and security domains. In Stackelberg games, one player, the leader, commits to a strategy and the follower makes her decision with knowledge of the leader's commitment. Existing algorithms for Stackelberg games efficiently find optimal solutions (leader strategy), but they critically assume that the follower plays optimally. Unfortunately, in many applications, agents face human followers (adversaries) who — because of their bounded rationality and limited observation of the leader strategy — may deviate from their expected optimal response. In other words, human adversaries' decisions are biased due to their bounded rationality and limited observations. Not taking into account these likely deviations when dealing with human adversaries may cause an unacceptable degradation in the leader's reward, particularly in security applications where these algorithms have seen deployment. The objective of this paper therefore is to investigate how to build algorithms for agent interactions with human adversaries.

To address this crucial problem, this paper introduces a new mixed-integer linear program (MILP) for Stackelberg games to consider human adversaries, incorporating: (i) novel *anchoring* theories on human perception of probability distributions and (ii) robustness approaches for MILPs to address human imprecision. Since this new approach considers human adversaries, traditional proofs of correctness or optimality are insufficient; instead, it is necessary to rely on empirical validation. To that end, this paper considers four settings based on real deployed security systems at Los Angeles International Airport (Pita et al., 2008 [35]), and compares 6 different approaches (three based on our new approach and three previous approaches), in 4 different observability conditions, involving 218 human subjects playing 2960 games in total. The final conclusion is that a model which incorporates both the ideas of robustness and anchoring achieves statistically significant higher rewards and also maintains equivalent or faster solution speeds compared to existing approaches.

* Corresponding author.
E-mail address: tambe@usc.edu (M. Tambe).

0004-3702/$ – see front matter
doi:10.1016/j.artint.2010.07.002

1. Introduction

In Stackelberg games, one player, the leader, commits to a strategy publicly before the remaining players, the followers, make their decision [16]. There are many multiagent security domains, such as attacker–defender scenarios and patrolling, where these types of commitments are necessary by the security agent [3,6,21,34] and it has been shown that Stackelberg games appropriately model these commitments [33,35]. For example, in an airport setting there may be six terminals serving passengers, but only four bomb sniffing canine units to patrol the terminals. In this scenario the canine units decide on a randomized patrolling strategy over these six terminals first, while their adversaries conduct surveillance and act taking this committed strategy into account. Indeed, Stackelberg games are at the heart of the ARMOR system deployed at the Los Angeles International Airport (LAX) to schedule security personnel since August 2007 [33,35] and have been deployed for the Federal Air Marshals service since October 2009 [21,45]. Moreover, these games have potential applications for network routing, pricing in transportation systems and many others [9,23].

Existing algorithms for Bayesian Stackelberg games find optimal solutions considering an *a priori* probability distribution over possible follower types [10,33]. Unfortunately, to guarantee optimality, these algorithms make strict assumptions on the underlying games, namely that the players are perfectly rational and that the followers perfectly observe the leader's strategy. However, these assumptions rarely hold in real-world domains, particularly when dealing with humans. Of specific interest are the security domains mentioned earlier (e.g. LAX) — even though an automated program may determine an optimal leader (security personnel) strategy, it must take into account a human follower (adversary). Such human adversaries may not be utility maximizers, computing optimal decisions. Instead, their decisions may be governed by their bounded rationality [41] which causes them to deviate from their expected optimal strategy. Humans may also suffer from limited observability of the security personnel's strategy, giving them a false impression of that strategy. In other words, when making decisions based on their own cognitive abilities, humans are biased due to their bounded rationality and inability to obtain complete sets of observations. Thus, a human adversary may not respond with the game theoretic optimal choice, causing the leader to face uncertainty over the gamut of adversary's actions. Therefore, in general, the leader in a Stackelberg game must commit to a strategy considering three different types of uncertainty:

(i) adversary response uncertainty due to her bounded rationality where the adversary may not choose the utility maximizing optimal strategy;
(ii) adversary response uncertainty due to her limitations in appropriately observing the leader strategy;
(iii) adversary reward uncertainty modeled as different reward matrices with a Bayesian *a priori* distribution assumption, i.e. a Bayesian Stackelberg game.

While existing algorithms handle the third type of uncertainty [10,33], these models can give a severely under-performing strategy when the adversary deviates because of the first two types of uncertainty. This degradation in leader rewards may be unacceptable in certain domains.

To overcome this limitation, this paper proposes a new algorithm based on a mixed-integer linear program (MILP). The major contribution of this new MILP is in providing a fundamentally novel integration of key ideas from:

(i) previous best known algorithms from the multiagent literature for solving Bayesian Stackelberg games;
(ii) robustness approaches for games from robust optimization literature [1,30];
(iii) anchoring theories on human perception of probability distributions from psychology [13,14,38].

While the robustness approach helps address human response imprecision, anchoring, which is an expansion of general support theory [46] on how humans attribute probabilities to a discrete set of events, helps address limited observational capabilities. To the best of our knowledge, the effectiveness of the combination of these ideas has not been explored in the context of Stackelberg games (or any other games). By uniquely incorporating these ideas our goal is to defend against the sub-optimal choices that humans may make due to bounded rationality or observational limitations. This MILP complements the prior algorithms for Bayesian Stackelberg games, handling all three types of uncertainty mentioned.[1]

Since this algorithm is centered on addressing non-optimal and uncertain human responses, traditional proofs of correctness and optimality are insufficient: it is necessary to experimentally test this new approach against existing approaches. Experimental analysis with human subjects allows us to show how this algorithm is expected to perform against human adversaries compared to previous approaches. To that end, we experimentally tested our new algorithm to determine its success by considering four settings based on real deployed security systems at LAX [35]. In all four settings, 6 different approaches were compared (three based on the new algorithm, one existing approach, and two baseline approaches), in 4 different observability conditions. These experiments involved 218 human subjects playing 2960 games in total and yielded statistically significant results showing that our new algorithm substantially outperformed existing methods when dealing with human adversaries. Runtime results were also gathered from our new algorithm against previous approaches showing that its solution speeds are equivalent to or faster than previous approaches. Based on these results we conclude that, while

[1] Although the MILP presented handles all three types of uncertainty, the focus of this paper is handling the first two types of uncertainty.