Adaptive and Dynamic Team Formation for Strategic and Tactical Planning

Sara Marie McCarthy
University of Southern California
sara.m.mccarthy@gmail.com

Introduction

Past work in security games has mainly focused on the problem static resource allocation; how to optimally deploy a given fixed team of resources. My research aims to address the challenge of integrating operational planning into security games, where resources are heterogeneous and the defender is tasked with optimizing over both the investment into these resources, as well as their deployment in the field. This allows the defender to design more adaptive strategies, reason about the efficiency of their use of these resources as well as their effectiveness in their deployment. This thesis explores the challenges in integrating these two optimization problems in both the single stage and multi-stage setting and provides a formal model of this problem, which we refer to as the Simultaneous Optimization of Resource Teams and Tactics (SORT) as a new fundamental research problem in security games that combines strategic and tactical decision making. The main contributions of this work are solution methods to the SORT problem under various settings as well as exploring various types of tradeoffs that can arise in these settings. These include managing budget for investment in resources as well as capacity constraints on use of resources. My work addresses scenarios when the tactical decision problem (optimal deployment) is difficult, and thus evaluating the performance of any given team is difficult. Additionally, I address domains where we are tasked with making repeated strategic level decision and where, due to changing domain features, fluctuations in time dependent processes or the realization of uncertain parameters in the problem, it becomes necessary to re-evaluate and adapt to new information.

Completed Work

Team Formation

I first provide a formal model of the SORT problem in the single stage setting (McCarthy et al. 2016b), where we are faced with two decision problems. Given a fixed budget to invest in a set of heterogeneous resources, each with different costs and capabilities, what is the optimal investment into these resources? And once the portfolio of security resources has been decided, how should those resources be combined into teams and deployed for optimal efficiency in the field? I studied this problem using the challenge of optimizing the defense of forests against illegal logging in Madagascar as a motivating domain, where I combine two different classes of security games (1) Green Security Games (Fang, Stone, and Tambe 2015) (Johnson et al. 2012), which deal with the protection of forests, fish and wildlife and (2) Network Security Games, where the challenge of simultaneously optimizing teams and deployment is particularly challenging due to the combinatorially large space of defender strategies (Jain et al. 2011). The fact that the tactical question is computationally challenging emphasizes the difficulty of the SORT problem, which requires evaluating the effectiveness of many teams to select the right one.The challenge of optimizing a team of security resources can be formulated as the following optimization problem:

$$\max_{\lambda \subset R} \left\{ F(\lambda) : \sum_{k \in \lambda} b_k \leq B \right\}$$

Where the value of a team of resources λ selected from some set of resources R is given by the expected utility of their optimal deployment, denoted $F(\lambda)$. $F(\lambda)$ can computationally difficult to calculate, particularly in Network Security Games, because it requires finding the optimal tactical allocation to assess the utility of a given team λ. Since there are an exponentially many possible teams, the sequential approach of evaluating $F(\lambda)$ exactly for every team and picking the best one is impractical. Instead, in my approach to SORT, I developed FORTIFY, a scalable branch and bound style algorithm which integrates the analysis of the strategic and tactical aspects of the problem to efficiently approximate $F(\lambda)$ and search the space of teams much more efficiently and limit the number of instances where we evaluate $F(\lambda)$ exactly. The novelty of FORTIFY is a hierarchical abstractions of the Network Security Game which provide bounds on $F(\lambda)$ the value of different teams of resources to speed up the search for the optimal team. I evaluated this work, using real network and resource data obtained from our partners working in Madagascar, and show that not only is the algorithm scalable, but as the number of possible combination of resources grows it because more and more worthwhile to perform this optimization over resources.

Active Sensing

I also looked at this problem in the multi-stage setting we are asked to make multiple investment decision, and where

the budget for investment in resources and cost of resources may change over time, due to time dependent features of the problem or information gain about uncertain problem parameters. One such problem is active sensing in a network. Here the defender may have many kinds of noisy security resources to monitor network traffic so each single alert does not provide a high confidence estimate about the security state (Sommer and Paxson 2010). Thus, the defender needs to come up with a sequential plan of actions, weighing the cost of deploying detectors to increase their knowledge and infer if an attack is taking place, with the potential loss due to successful attacks as well as the cost of misclassifying legitimate network use, and determine the best response policy. To address this, I developed a novel decision-theoretic model Virtually Distributed Partially Observable Markov Decision Process (VD-POMDP) to reason about noisy observations in order to dynamically allocate resources (McCarthy et al. 2016a). The efficiency of the formulation is based on two key contributions: (i) the problem is decomposed in a way that allows for individual sub-POMDPs with sparse interactions to be created. Individual policies for different sub-POMDPs are planned separately and their sparse interactions are only resolved at execution time to determine the joint actions to perform; (ii) The abstraction in planning step allows for large speedups and scalability, after which a fast MILP is used to implement the abstraction while resolving any interactions. I provide conditions under which these methods result in an optimal joint policy, and provide empirical evidence high solution quality when these conditions are not satisfied. I also provide experimental evaluation of our model in a real network testbed, where I demonstrate the ability to correctly identify real attacks.

Adaptive Screening

Here I consider the problem of dynamically allocating screening resources of different efficacies (e.g., magnetic or X-ray imaging) at checkpoints (e.g., at airports) to successfully avert an attack by one of the screenees. Previously, the Threat Screening Game model (Brown et al. 2016) was introduced to address this problem under the assumption that screenee arrival times are perfectly known. In reality, arrival times are uncertain, which severely impedes the implementability and performance of this approach. When these uncertain parameters are realized in a way that is different that what had been planned for this can result in not only sub-optimal screening strategies, but also in strategies that are not implementable. The original TSG formulation fails in its mission to realistically model real-world settings. Addressing this challenge is difficult, as it requires reasoning about all the possible realizations of the uncertainty and coming up with an optimal plan for each of those scenarios, one which appropriately balances the tradeoffs between throughput efficiency and screening effectiveness. To address this, I introduce a novel framework for dynamic allocation of threat screening resources known as Robust Threat Screening Games (RTSG) that explicitly accounts for uncertainty in the screenee arrival times, as well as the coupling of capacity constraints in time(McCarthy, Vayanos, and Tambe 2017). I provide a tractable solution approach using compact linear decision rules combined with robust reformulation and constraint randomization. I performed extensive experiments which showcase that my approach outperforms *(i)* exact solution methods in terms of tractability, while incurring only a very minor loss in optimality, and *(ii)* methods that ignore uncertainty in terms of both feasibility and optimality.

Future Work

Looking forward, I am currently interested in exploring problem domains where AI may be used to address challenges in social welfare and social good. I plan to expand on my work in the domain of green security and look at improving the scalability of my solution methods to the SORT problem, addressing the challenging case where it is infeasible to consider the entire space of team which may be formed. I also plan on developing solution techniques that are robust to uncertainty in the strategic planning problem. It may be the case that there is uncertainty in the availability of different resources, their costs or in the budget available for investment. This is particularly relevant in the multistage setting where we need to make decisions now based uncertain information about how we may invest in the future.

References

Brown, M.; Sinha, A.; Schlenker, A.; and Tambe, M. 2016. One size does not fit all: A game-theoretic approach for dynamically and effectively screening for threats. In *AAAI conference on Artificial Intelligence (AAAI)*.

Fang, F.; Stone, P.; and Tambe, M. 2015. When security games go green: Designing defender strategies to prevent poaching and illegal fishing. In *International Joint Conference on Artificial Intelligence (IJCAI)*.

Jain, M.; Korzhyk, D.; Vaněk, O.; Conitzer, V.; Pěchouček, M.; and Tambe, M. 2011. A double oracle algorithm for zero-sum security games on graphs. In *AAMAS*, 327–334.

Johnson, M. P.; Fang, F.; ; and Tambe, M. 2012. Patrol strategies to maximize pristine forest area. In *Conference on Artificial Intelligence (AAAI)*.

McCarthy, S.; Sinha, A.; Tambe, M.; and Manadhata, P. 2016a. Data exfiltration detection and prevention: Virtually distributed pomdps for practically safer networks. In *Decision and Game Theory for Security (GameSec 2016)*.

McCarthy, S.; Tambe, M.; Kiekintveld, C.; Gore, M. L.; and Killion, A. 2016b. Preventing illegal logging: Simultaneous optimization of resource teams and tactics for security. In *AAAI conference on Artificial Intelligence (AAAI)*.

McCarthy, S.; Vayanos, P.; and Tambe, M. 2017. Staying ahead of the game: Adaptive robust optimization for dynamic allocation of threat screening resources. In *International Joint Conference on Artificial Intelligence (IJCAI)*.

Sommer, R., and Paxson, V. 2010. Outside the closed world: On using machine learning for network intrusion detection. In *Security and Privacy (SP), 2010 IEEE Symposium on*, 305–316. IEEE.