Extended White Paper

Game Theory for Security: Large-scale data-driven approaches

Modern security problems are rapidly growing in their complexity and scale, while the potential severity of losses grows in parallel. Real-world security agencies have scarce resources with which they are required to protect against intelligent and poorly understood adversaries. They need to deploy these resources effectively. Game theory is already a well-established tool in security problems, since it is the study of conflict. Robust optimization offers tools to make decisions in the presence of uncertainty, and it has experienced considerable success in Operations Research. Machine learning and statistical learning allow us to make use of data as it becomes available and to update our beliefs about the situation at hand. Taken together, game theory, robust optimization, and learning are a powerful triumvirate.

It is our philosophy to have only real-world problems drive our research, and we measure our success in terms of practical deployment of our work. It is this attention to real-world problems that has enhanced our interest in further coupling robust optimization and learning with game theory, since real-world security games are played under considerable uncertainty, yet data is available that can aid our planning.

Project Technical Description

1. Theme Area: Risk Analysis – model and mitigate the risk of losses from the attacks of intelligent, adaptive adversaries.

2. Principal Investigator:
Milind Tambe
3. Institution: USC
4. **Co-Investigators:** None
5. **Keywords:** security games, risk management, robust optimization, learning, mathematical programming, large-scale implementation, predictive analytics

New domains for game theory: A TRUSTS-scheduled patrol to check fares on METRO trains in Los Angeles; a TRUSTS schedule for LA Sheriff's Dept officers to patrol METRO trains shown on a handheld device; and patrol for protection of fisheries (picture shows a Mexican lancha seized by the US Coast Guard).

6. Brief Description:

Game-theoretic methods have been the central theme of our work on real-world security problems. In particular, we have emphasized Stackelberg games, games between a 'defender' and an 'attacker'. In our setting the defender is the security agency, while the attacker is an aggressor who can observe the leader's play and then choose where and how to attack. Our research has led to many successful deployments, with measureable improvements in safety. IRIS has been deployed by the Federal Air Marshals Service (FAMS) since 2009. PROTECT has been deployed by the US Coast Guard in Boston since 2011 and New York since 2012, and it is slated for nationwide deployment. Finally, GUARDS is being evaluated by the TSA for infrastructure protection.

Currently, TRUSTS is being evaluated by the Los Angeles Sheriff's Dept (LASD) on the Los Angeles Metro Rail system. TRUSTS is a data-driven system that intelligently schedules patrols. We have built hand-held mobile device interfaces to coordinate patrols, as well as to collect large

amounts of data. The scale of this problem is very large, and there is an abundance of data, these two features make it difficult to use existing methods because they do not scale to our setting. In another stream of current research, we are aiding the US Coast Guard in their mission to protect fisheries in the Gulf of Mexico from Mexican lanchas (illegal fisherman). In this work, we are given large amounts of data about lancha sightings, Coast Guard patrols, fish stocks, and weather, and we must extract meaning from these data to understand adversary movement and behavior. Additionally, we want to use what we learn from the data to design better and more robust Coast Guard patrols. In parallel work, we are studying the problem of wildlife protection, which also leads to large-scale data-driven problems.

In our methodological research, we have begun applying robust optimization to develop robust Stackelberg security games. Our formulations cover a wide-range of possible model uncertainty, and the resulting solutions to these robust Stackelberg games offer greater reliability to the defender than classical Stackelberg games. We are currently extending this work to allow for different levels of optimism/pessimism in the design of uncertainty sets, and to incorporate risk-awareness.

Our recent investigations in both theory and practice have led us to a new class of games beyond our original work on Stackelberg games. This new class of games is large-scale, dynamic, overflowing with data, and includes multiple heterogeneous defender assets. These existing features push this new class beyond the current theory. Furthermore, we face model ambiguity about the nature and the objectives of the attacker. Existing computational methods cannot handle games of this class, so we are pushing the theoretical frontier to make these models practical.

7. Objectives:

Our proposed research will address the following fundamental research challenges. Though we describe each of these challenges in isolation, this project will produce a comprehensive theory that addresses all of these research challenges simultaneously.

(a) *Large-scale Implementation:* We have successfully applied our work to many real-world security games in the past, the issue of scaling up our methods to multi-stage games has not been resolved. In particular, both the LA Sheriff's Dept and US Coast Guard are looking for patrols over long time horizons, over large areas.

(b) *Learning and Data-Driven Optimization:* We are exploring new domains on combating *everyday crimes*, this type of crime differs from the strategic crimes faced in counter-terrorism applications. A terrorist attack is carefully planned and implemented, while everyday crimes are unplanned and opportunistic. Our work on the protection of fishery and marine resources, mass transit systems such as trains and roads, and crime hot spots around cities, all falls under the heading of opportunistic crime. In cooperation with law enforcement, it is possible to collect abundant data on this class of security problem. In current research, we have developed a Bayesian game representation of the LASD's metro problem that allows for learning about the underlying parameters of the game. In our proposed work, we will incorporate online learning into many other security games, and search for data-driven strategies.

(c) *Robust optimization:* Behavioral game theory experiments have revealed that players often deviate from perfectly rational play. While we can use data and repeated interaction to learn about our adversaries, in life-critical security games we need strong robustness guarantees on our policies. Certainly, one could argue that many terrorists are not

perfectly rational decision makers because of the high cost of their attacks to themselves. So, we need strong and tractable models that can handle uncertainty in attacker behavior. Our work on robust Stackelberg games is promising and we can extend these methods to incorporate multiple types of uncertainty in a dynamic setting.

(d) *Coordination:* In many cases we are faced with highly heterogeneous defender units, potentially belonging to different agencies. For example, in New York we may potentially need to coordinate US Coast Guard boats, helicopter patrols, and NYPD boat patrols. Most of the patrol problems we have discussed so far are very large-scale and difficult to solve for even a single patroller unit. To effectively coordinate between patrol units, we need a theory for decomposing the game into computationally manageable pieces, in a way that still gives strong performance guarantees.

8. Interfaces to CREATE Projects:

This work will maintain a close interface with CREATE's PortSec risk analysis and economics project in evaluating the security resource allocations, and CREATE's adaptive adversary project. In both cases, our collaboration will emphasize large-scale data-driven applications.

9. Previous or current work relevant to the proposed project:

Above, we have discussed several previous applications and relevant research. While our previous work has been on developing game theoretic applications, each one has surpassed the last in terms of complexity of models of adversary behavior, scale of the problem, types and numbers of adversaries, and constraints accommodated. Our newer domains, TRUSTS and US fishery protection, are more dynamic domains, with large amounts of data, that requie coordination of multiple defender resources. These new domains have lead to significant new research challenges. Our resulting solutions are certain to have wide impact across the field.

10. Major Deliverables, Research Transition Products and Customers:

- **Products:** Our previous research has resulted in several deployed software applications for homeland security agencies. Our current research will focus on the TRUSTS software for the LASD and new software for the US Coast Guard for protection of fish. As opportunities develop to liaise with Customs and Border Protection and other agencies, our research will also focus on those opportunities. Our newer software will be heavily data-driven.
- **Customers:** The FAMS service has been using IRIS since 2009; PROTECT has been in use by the US Coast Guard since 2011 in Boston and since 2012 in New York, it will be deployed nationwide; GUARDS is under evaluation by the TSA. Our new software will be usable by the Los Angeles Sheriff's department (TRUSTS) and the US Coast Guard (fisheries).

11. Technical Approach:

A more detailed plan for our research follows:

(e) *Large-scale Implementation:* For the problem of dynamic patrol scheduling, we are taking an approach based on Markov Decision Processes (MDPs). These methods necessarily suffer from the curse of dimensionality because of their scale. In response, we are exploring large-scale linear programming methods, such as duality decomposition, as we have had success using double oracle methods in earlier work. We are also interested in decomposition schemes for dynamic stochastic games more generally. Given an appropriate theory of decomposition, we can justify breaking large games apart into

smaller, more tractable ones, and then we can recover a good strategy for the original game. **Simulation-based algorithms**, such as Q-learning and optimistic policy iteration, have been very successful in the solution of large-scale problems in optimization and control. There is also initial theoretical work on the application of Q-learning in Nash games. It is a very natural idea to extend these simulation-based algorithms to the multi-agent setting. **Decomposition** is another highly successful tool from optimization and control that is designed to facilitate large-scale implementation. It is a very natural idea to intelligently break apart large-games into smaller ones to make their solution easier.

a) *Learning and Data-Driven Optimization:* We will investigate new learning approaches for the everyday-crime domains, by combining game theory and machine learning techniques. In particular, for the fishery-protection domain, the defender (Coast Guard) may be uncertain about the locations of the fish stock, whereas the adversary (illegal fishermen) may have more information. Nevertheless, data from repeated interactions allow us to learn the locations of fish stock by observation. We will use the latest hot spot detection models to help us track both the fish and the lanchas, and then we will compute robust patrol routes given our predictions of where the lanchas will go. More generally, learning in games alludes to the exploration versus exploration tradeoff, common in all learning problems. We should try new things because they might work well (exploration), but if we have something that works well we should be conservative and keep using it (exploitation). A careful balance between these two elements is needed, especially in the life critical systems that we study.
b) *Robust optimization:* Behavioral game theory and psychology offer various models to explain human decision-making in games. The Quantal Response model has provided significant advantages over the perfect rationality model, but it leads to non-convex and non-linear optimization problem, which are difficult to solve. We are interested in robust models of adversary behavior for exactly this reason. Not only do many robust optimization formulations enjoy tractability, but they offer strong performance guarantees. Our earlier work has leveraged this success for adversary uncertainty in Stackelberg games. Our earlier work emphasized a specific type of uncertainty set which led to highly conservative and easily computable strategies. Now, we are considering more general uncertainty sets which allow for greater optimism on the part of the defender. The defender wants to put the most effective security system in practice possible, but it would be sub-optimal to spend far more than is necessary to do so. We need more flexible uncertainty sets that do not always drive the decision maker to plan against the absolute worst-case. The next step is to combine robustness with learning to create robust learning algorithms.
c) *Coordination:* Coordination between patrol assets is difficult because of the cross effects of joint actions, i.e. two patrol units defending nearby resources may have greater effectiveness than the sum of their individual effectiveness. This type of interaction creates computational challenges for scaling up to large numbers of units, because we cannot simply break apart the game and solve it separately for each unit. In one approach, we plan to exploit the submodular structure of utility functions and apply techniques from submodular optimization to yield approximate algorithms with theoretical guarantees on solution quality. We also plan to adapt models and techniques from multiagent cooperative planning, including the Decentralized Markov Decision Process (Dec-MDP) model. This problem falls under the more general heading of decomposition techniques, so we will bring this toolset to bear here as well.

12. References

1. M. Jain, B. An, M. Tambe: An Overview of Recent Application Trends at the AAMAS conference: Security, Sustainability and Safety. AI Magazine, 2012
2. F. Ordonez, M. Tambe, J. F. Jara, M. Jain, C. Kiekintveld, J. Tsai: Deployed Security Games for Patrol Planning, Handbook on Operations Research for Homeland Security (Book chapter) 2012
3. E. Shieh, B. An, R. Yang, M. Tambe, C. Baldwin, J. DiRenzo, B. Maule, G. Meyer: PROTECT: A Deployed Game Theoretic System to Protect the Ports of the United States. International Conference on Autonomous Agents and Multiagent Systems (AAMAS) 2012
4. Z. Yin, M. Tambe: A Unified Method for Handling Discrete and Continuous Uncertainty in Bayesian Stackelberg Games, International Conference on Autonomous Agents and Multiagent Systems (AAMAS) 2012
5. R. Yang, F. Ordonez, M. Tambe: Computing Optimal Strategy against Quantal Response in Security Games, International Conference on Autonomous Agents and Multiagent Systems (AAMAS) 2012
6. Y. Vorobeychik, B. An, M. Tambe: Adversarial Patrolling Games: Extended Abstract. International Conference on Autonomous Agents and Multiagent Systems (AAMAS) (Short paper) 2012
7. J. Pita, R. John, R. Maheswaran, M. Tambe, R. Yang, S. Kraus: A Robust Approach to Addressing Human Adversaries in Security Games: Extended Abstract. International Conference on Autonomous Agents and Multiagent Systems (AAMAS) (Short paper) 2012
8. TRUSTS: Scheduling Randomized Patrols for Fare Inspection in Transit Systems. Zhengyu Yin, Albert Jiang, Matthew Johnson, Milind Tambe, Christopher Kiekintveld, Kevin Leyton-Brown, Tuomas Sandholm, John Sullivan. Conference on Innovative Applications of Artificial Intelligence (IAAI) 2012
9. R. Yang, F. Fang, A. X. Jiang, K. Rajagopal, M. Tambe, R. Maheswaran: Designing Better Strategies against Human Adversaries in Network Security Games: Extended Abstract. International Conference on Autonomous Agents and Multiagent Systems (AAMAS)(Short paper) 2012
10. Security Games with Limited Surveillance. Bo An, David Kempe, Christopher Kiekintveld, Eric Shieh, Satinder Singh, Milind Tambe, Yevgeniy Vorobeychik. Conference on Artificial Intelligence (AAAI) 2012
11. PROTECT: An Application of Computational Game Theory for the Security of the Ports of the United States. Eric Shieh, Bo An, Rong Yang, Milind Tambe, Craig Baldwin, Joseph DiRenzo, Ben Maule, Garrett Meyer. Conference on Artificial Intelligence (AAAI) Spotlight Track 2012
12. Security Games for Controlling Contagion. Jason Tsai, Thanh H. Nguyen, Milind Tambe. Conference on Artificial Intelligence (AAAI) 2012
13. The Deployment-to-Saturation Ratio in Security Games. Manish Jain, Kevin Leyton-Brown, Milind Tambe. Conference on Artificial Intelligence (AAAI) 2012
14. Patrol Strategies to Maximize Pristine Forest Area. Matthew P. Johnson, Fei Fang, and Milind Tambe. Conference on Artificial Intelligence (AAAI) 2012
15. A Robust Approach to Addressing Human Adversaries in Security Games. James Pita, Richard John, Rajiv Maheswaran, Milind Tambe, Sarit Kraus. European Conference on Artificial Intelligence (ECAI) 2012
16. Analysis of Heuristic Techniques for Controlling Contagion. Jason Tsai, Nicholas Weller, Milind Tambe. AAAI Fall Symposium, 2012
17. Security Games on Social Networks. Thanh H. Nguyen, Jason Tsai, Albert Jiang, Emma Bowring, Rajiv Maheswaran, Milind Tambe. AAAI Fall Symposium, 2012.
18. Game Theory for Security: Key Algorithmic Principles, Deployed Systems, Lessons Learned. Milind Tambe, Manish Jain, James Adam Pita, Albert Xin Jiang. 50th Annual Allerton Conference on Communication, Control, and Computing, 2012
19. A Deployed Quantal Response Based Patrol Planning System for the US Coast Guard. Bo An, Fernando Ordonez, Milind Tambe, Eric Shieh, Rong Yang, Craig Baldwin, Joseph DiRenzo, Ben Maule, Garrett Meyer. Interfaces 2013

20. Improving Resource Allocation Strategies Against Human Adversaries in Security Games: An Extended Study. R. Yang, C. Kiekintvled, F. Ordonez, M. Tambe, R. John. Artificial Intelligence Journal (AIJ), (to appear) 2013
21. Empirical Evaluation of Computational Fear Contagion Models in Crowd Dispersions. Jason Tsai, Emma Bowring, Stacy Marsella, Milind Tambe. Journal of Autonomous Agents and Multiagent Systems, JAAMAS (to appear) 2013

13. Major Milestones and Dates:

1. Develop faster algorithms games with network representations and other cases — Ongoing through Dec 2013
2. Conduct behavioral experiments, perform data analysis to generate improved models of human adversaries — Ongoing through Dec 2013
3. New solution methods based on behavioral models — March 2014
4. Exploration and development of machine learning integration with game theory — Sept 2014
5. Coordination of multiple defender units — March 2014
6. Teach short course/executive education — TBD