

REPORT DOCUMENTATION PAGE			Form Approved OMB NO. 0704-0188		
The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA, 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YYYY) 15-08-2016		2. REPORT TYPE Final Report		3. DATES COVERED (From - To) 15-Aug-2011 - 14-May-2015	
4. TITLE AND SUBTITLE Final Report: Data-Driven Game-Theory			5a. CONTRACT NUMBER W911NF-11-1-0344		
			5b. GRANT NUMBER		
			5c. PROGRAM ELEMENT NUMBER 1M30BM		
6. AUTHORS V.S. Subrahmanian			5d. PROJECT NUMBER		
			5e. TASK NUMBER		
			5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAMES AND ADDRESSES University of Maryland - College Park Office of Research Administration 3112 Lee Building 7809 Regents Drive College Park, MD 20742 -5141			8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS (ES) U.S. Army Research Office P.O. Box 12211 Research Triangle Park, NC 27709-2211			10. SPONSOR/MONITOR'S ACRONYM(S) ARO		
			11. SPONSOR/MONITOR'S REPORT NUMBER(S) 59726-NS.56		
12. DISTRIBUTION AVAILABILITY STATEMENT Approved for Public Release; Distribution Unlimited					
13. SUPPLEMENTARY NOTES The views, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy or decision, unless so designated by other documentation.					
14. ABSTRACT We make multiple contributions in this research. Though classical game theory assumes that payoff matrices are given for the players involved in the game, we show how these can be learned automatically from a body of data. We further show new definitions of quasi-polynomially computable approximate equilibria and show how to efficiently and approximately compute them. We apply these methods to the behavior of real world terrorist groups. Continuing with our work on adversarial models, we study methods to disclose information publicly in order to share the adversary's actions to our advantage. We further develop methods to quantify and reduce lethality of					
15. SUBJECT TERMS Game Theory, Reasoning, Geospatial					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	15. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON Venkatramanan Subrahmanian
a. REPORT UU	b. ABSTRACT UU	c. THIS PAGE UU			19b. TELEPHONE NUMBER 301-405-6724

Report Title

Final Report: Data-Driven Game-Theory

ABSTRACT

We make multiple contributions in this research. Though classical game theory assumes that payoff matrices are given for the players involved in the game, we show how these can be learned automatically from a body of data. We further show new definitions of quasi-polynomially computable approximate equilibria and show how to efficiently and approximately compute them. We apply these methods to the behavior of real world terrorist groups. Continuing with our work on adversarial models, we study methods to disclose information publicly in order to shape the adversary's actions to our advantage. We further develop methods to quantify and reduce lethality of networks. Our work develops and further studies the use of game-theoretic frameworks to the study of security games, online games, as well as games involving diverse terrorist groups including Lashkar-e-Taiba and the Indian Mujahideen.

Enter List of papers submitted or published that acknowledge ARO support from the start of the project to the date of this printing. List the papers, including journal references, in the following categories:

(a) Papers published in peer-reviewed journals (N/A for none)

Received

Paper

01/06/2016 36.00 Ron Katz, Sarit Kraus, Rina Azoulay. Efficient bidding strategies for Cliff-Edge problems, Autonomous Agents and Multi-Agent Systems, (04 2013): 290. doi:

01/06/2016 40.00 Francesca Spezzano, V.S. Subrahmanian, Aaron Mannes. Reshaping Terrorist Networks, Communications of the ACM, (08 2014): 60. doi:

01/06/2016 41.00 Edoardo Serra, V.S. Subrahmanian. A Survey of Quantitative Models of Terror Group Behavior and an Analysis of Strategic Disclosure of Behavioral Models, IEEE Transactions on Computational Social Systems, (03 2014): 66. doi:

TOTAL: 3

Number of Papers published in peer-reviewed journals:

(b) Papers published in non-peer-reviewed journals (N/A for none)

Received

Paper

TOTAL:

Number of Papers published in non peer-reviewed journals:

(c) Presentations

Non Peer-Reviewed Conference Proceeding publications (other than abstracts):

<u>Received</u>	<u>Paper</u>
08/15/2016 1.00	Galit Haim, Ya'akov (Kobi) Gal, Sarit Kraus, Michele Gelfand. Learning to Reveal Information in Repeated Human-Computer Negotiation, Human-Agent Interaction Design and Models, HAIDM 2012 Workshop. 04-JUN-12, Valencia, Spain. : ,
08/15/2016 51.00	. Model-Predictive Target Defense by Team of Unmanned Surface Vehicles Operating in Uncertain Environments, International Conference on Robotics and Automation ICRA 2013. 06-MAY-13, Karlsruhe, Germany. : ,
08/15/2016 49.00	. Strategic Advice Provision in Repeated Human-Agent Interactions, AAAI 2012. 22-JUL-12, Toronto, Ontario, Canada. : ,
08/15/2016 55.00	. A New Paradigm for the Study of Corruption in Different Cultures, SBP14. 01-APR-14, Washington, D.C.. : ,
08/15/2016 48.00	. Automated Strategies for Determining Rewards for Human Work, 26th AAAI Conference on Artificial Intelligence. 22-JUL-12, Toronto, Ontario, Canada. : ,
08/15/2016 24.00	Moshe Bitan, Ya'akov Gal, Sarit Kraus, Elad Dokow, Amos Azaria. Social Rankings in Human-Computer Committees, 27th AAAI Conference on Artificial Intelligence. 14-JUL-13, Bellevue, WA. : ,
08/15/2016 21.00	Noam Peled, Moshe Bitan, Joseph Keshet, Sarit Kraus. Predicting Human Strategic Decisions Using Facial Expression, International Joint Conference on Artificial Intelligence. 03-AUG-13, Beijing, China. : ,
08/15/2016 30.00	Amos Azaria, Sarit Kraus. Advice Provision in Multiple Prospect Selection Problems, HAIDM 2013. 06-MAY-13, Saint Paul, MN. : ,
08/15/2016 28.00	Amos Azaria, Ariella Richardson, Sarit Kraus. Autonomous Agent for Deception Detection, HAIDM 2013. 06-MAY-13, Saint Paul, MN. : ,
08/15/2016 25.00	Thanh H. Nguyen, Rong Yang, Amos Azaria, Sarit Kraus, Milind Tambe. Analyzing the Effectiveness of Adversary Modeling in Security Games, AAAI Conference on Artificial Intelligence 2013. 14-JUL-13, Bellevue, WA. : ,
08/15/2016 18.00	John P. Dickerson, Anshul Sawant, Mohammad T. Hajiaghayi, V.S. Subrahmanian. PREVE: A Policy Recommendation Engine based on Vector Equilibria Applied to Reducing LeT's Attacks, ASONAM 2013. 26-AUG-13, Niagara Falls, Canada. : ,
08/15/2016 5.00	Asaf Frieder, Raz Lin, Sarit Kraus. Agent-human Coordination with Communication Costs under Uncertainty (Extended Abstract), Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems 2012. 04-JUN-12, Valencia, Spain. : ,
08/15/2016 3.00	Noam Peled, Kobi Gal, Sarit Kraus. Learning to Reveal Information in Repeated Human-Computer Negotiation, Human-Agent Interaction Design and Models, HAIDM 2012 Workshop. 04-JUN-12, Valencia, Spain. : ,

08/26/2013	26.00	Albert Xin Jiang, Zhengyu Yin, Chao Zhang, Milind Tambe, Sarit Kraus. Game-theoretic Randomization for Security Patrolling with Dynamic Execution Uncertainty, Autonomous Agents and Multiagent Systems. 06-MAY-13, . : ,
08/26/2013	23.00	Noam Peled, Ya'akov (Kobi) Gal, Sarit Kraus. An Agent Design for Repeated Negotiation and Information Revelation with People, AAAI Conference on Artificial Intelligence. 14-JUL-13, . : ,
08/26/2013	22.00	Samuel Barrett, Peter Stone, Sarit Kraus, Avi Rosenfeld. Teamwork with Limited Knowledge of Teammates, AAAI Conference on Artificial Intelligence. 14-JUL-13, . : ,
08/26/2013	29.00	Amos Azaria, Ariella Richardson, Avshalom Elmalech, Avi Rosenfeld, Sarit Kraus, David Sarne. On Automated Agents' Rationality, HAIDM 2013. 06-MAY-13, . : ,
10/04/2012	4.00	Moshe Bitan, Ya'akov (Kobi) Gal, Elad Dokaw, Sarit Kraus. Social Rankings in Human-Computer Committees, Human-Agent Interaction Design and Models, HAIDM 2012 Workshop. 04-JUN-12, . : ,

TOTAL: 18

Number of Non Peer-Reviewed Conference Proceeding publications (other than abstracts):

Peer-Reviewed Conference Proceeding publications (other than abstracts):

<u>Received</u>	<u>Paper</u>
08/15/2016	8.00 Brandon Wilson, Inon Zuckerman, Austin Parker, Dana S. Nau. Improving Local Decisions in Adversarial Search, European Conference on Artificial Intelligence 2012. 27-AUG-12, Montpellier, France. : ,
08/15/2016	38.00 Amos Azaria, Ariella Richardson, Sarit Kraus. An Agent for Deception Detection in Discussion Based Environments, 18th ACM Conference on Computer-Supported Cooperative Work and Social Computing. 14-MAR-14, Vancouver, Canada. : ,
08/15/2016	35.00 Avi Rosenfeld, Inon Zuckerman, Erel Segal-Halevi, Osnat Drein, Sarit Kraus. NegoChat: A Chat-Based Negotiation Agent, International Conference on Autonomous Agents and Multiagent Systems. 05-MAY-14, Paris, France. : ,
10/13/2014	33.00 Kan-Leung Cheng, Inon Zuckerman, Dana Nau, Jennifer Golbeck. Predicting Agents' Behavior by Measuring their Social Preferences, European Conference on Artificial Intelligence 2014. 18-AUG-14, . : ,

TOTAL: 4

(d) Manuscripts

<u>Received</u>	<u>Paper</u>
01/12/2015 43.00	Patrick Roos, Michele Gelfand, Dana Nau, Janetta Lun. Societal Threat and Cultural Variation in the Strength of Social Norms:An Evolutionary Basis, Organizational Behavior and Human Decision Processes (12 2014)
07/10/2013 13.00	Brandon Wilson, Inon Zuckerman, Austin Parker, Dana S. Nau. Improving Local Decisions in Adversarial Search, TBD (06 2013)
07/10/2013 12.00	Eric Raboin, Petr Svec, Dana Nau, Satyandra K. Gupta. Model-Predictive Target Defense by Team of Unmanned Surface Vehicles Operating in Uncertain Environments, TBD (06 2013)
08/15/2016 2.00	Amos Azaria, Yonatan Aumann, Sarit Kraus. Automated Strategies for Determining Rewards for Human Work, Association for the Advancement of Artificial Intelligence (08 2012)
08/15/2016 39.00	Amos Azaria, Yonatan Aumann, Sarit Kraus. Automated Agents for Reward Determination for Human Work in Crowdsourcing Applications, Journal of Autonomous Agents and Multiagent System (09 2014)
08/15/2016 37.00	Ya'akov Gal, Avi Rosenfeld, Sarit Kraus, Michele Gelfand, Bo An, Jun Lin. A New Paradigm for the Study of Corruption in Different Cultures, Proc. of 2014 International Conference on Social Computing, Behavioral-Cultural Modeling & Prediction (09 2014)
08/15/2016 32.00	Anshul Sawant, John P. Dickerson, Mohammad T. Hajiaghayi, V.S. Subrahmanian. Automated Generation of Counter-Terrorism Policies using Multi-Expert Input, ACM Transactions on Intelligent Systems and Technology (09 2014)
08/15/2016 27.00	Avi Rosenfeld, Zevi Bareket, Claudia V. Goldman, Sarit Kraus, David J. LeBlanc, Omer Tsimhoni. Toward Adapting Cars to Their Drivers, AI Magazine (10 2012)
08/15/2016 16.00	Cristian Molinaro, Amy Sliva, V.S. Subrahmanian. Super-Solutions: Succinctly Representing Solutions in Abductive Annotated Probabilistic Temporal Logic, ACM Transactions of Computational Logic (07 2013)
08/15/2016 11.00	Dana S. Nau, ?. Theoretical Analysis and Evolutionary Simulation of Social-Learning Strategies, Journal of Artificial Societies and Social Simulation (JASSS) (06 2013)
08/15/2016 6.00	Amos Azaria, Zinovi Rabinovich, Sarit Kraus, Claudia V. Goldman, Ya'akov Gal. Strategic Advice Provision in Repeated Human-Agent Interactions, Association for the Advancement of Artificial Intelligence (07 2012)
08/15/2016 7.00	Avi Rosenfeld, Inon Zukerman, Amos Azaria, Sarit Kraus. Combining Psychological Models with Machine Learning, Synthese (10 2012)
08/16/2016 9.00	Ryan Carr, Eric Raboin, Austin Parker, Dana Nau. Theoretical and Experimental Analysis of an Evolutionary Social-Learning Game, Submitted for publication (10 2012)

08/23/2013 17.00 Francesco Parisi, Amy Sliva, V.S. Subrahmanian. Embedding Forecast Operators in Databases, Elsevier (02 2012)

08/26/2013 31.00 Meirav Hadad, Sarit Kraus, Irth Ben-Arroyo Hartman, Avi Rosenfeld. Group planning with time constratins, Annals of Mathematics and Artificial Intelligence (06 2013)

10/13/2014 34.00 Patrick Roos, Michele Gelfand, Dana Nau, Ryan Carr. High strength-of-ties and low mobility enable the evolution of third-party punishment, Proceedings of the Royal Society B: Biological Sciences (12 2013)

10/13/2014 42.00 Edoardo Serra, Francesca Spezzano, V.S. Subrahmanian. ChoiceGAPs: Competitive Diffusion as a Massive Multi-Player Game in Social Networks, ACM Transactions of Computational Logic (08 2014)

TOTAL: 17

Number of Manuscripts:

Books	
<u>Received</u>	<u>Book</u>
TOTAL:	

Received Book Chapter

08/15/2016 47.00 . Region-Based Geospatial Abduction with Counter-IED Applications, : Springer-Verlag, (2011)

TOTAL: 1

Patents Submitted

Patents Awarded

Awards

Graduate Students

<u>NAME</u>	<u>PERCENT SUPPORTED</u>	Discipline
Anshul Sawant	0.00	
Manish Purohit	0.10	
Brittany Schuetzle	0.00	
Paulo Shakarian	0.20	
Brandon Wilson	0.10	
Ryan Carr	0.20	
Eric Raboin	0.20	
Amy Sliva	0.10	
Patrick Roos	0.50	
FTE Equivalent:	1.40	
Total Number:	9	

Names of Post Doctorates

<u>NAME</u>	<u>PERCENT SUPPORTED</u>
Aaron Mannes	0.80
Edoardo Serra	0.50
Francesca Spezzano	0.40
FTE Equivalent:	1.70
Total Number:	3

Names of Faculty Supported

<u>NAME</u>	<u>PERCENT SUPPORTED</u>	National Academy Member
V.S. Subrahmanian	0.05	
Dana S. Nau	0.07	
FTE Equivalent:	0.12	
Total Number:	2	

Names of Under Graduate students supported

<u>NAME</u>	<u>PERCENT SUPPORTED</u>
FTE Equivalent:	
Total Number:	

Student Metrics

This section only applies to graduating undergraduates supported by this agreement in this reporting period

The number of undergraduates funded by this agreement who graduated during this period: 0.00

The number of undergraduates funded by this agreement who graduated during this period with a degree in science, mathematics, engineering, or technology fields:..... 0.00

The number of undergraduates funded by your agreement who graduated during this period and will continue to pursue a graduate or Ph.D. degree in science, mathematics, engineering, or technology fields:..... 0.00

Number of graduating undergraduates who achieved a 3.5 GPA to 4.0 (4.0 max scale):..... 0.00

Number of graduating undergraduates funded by a DoD funded Center of Excellence grant for Education, Research and Engineering:..... 0.00

The number of undergraduates funded by your agreement who graduated during this period and intend to work for the Department of Defense 0.00

The number of undergraduates funded by your agreement who graduated during this period and will receive scholarships or fellowships for further studies in science, mathematics, engineering or technology fields: 0.00

Names of Personnel receiving masters degrees

NAME

Total Number:

Names of personnel receiving PHDs

NAME

Total Number:

Names of other research staff

NAME

PERCENT SUPPORTED

FTE Equivalent:

Total Number:

Sub Contractors (DD882)

Inventions (DD882)

Scientific Progress

Technology Transfer

ARO Grant ARO Grant W911NF1110344

Report submitted to ARO for the period April 2011 – August 2015

V.S. Subrahmanian

University of Maryland College Park

College Park, MD 20742.

vs@cs.umd.edu

301-405-6724

Abstract. We make multiple contributions in this research. Though classical game theory assumes that payoff matrices are given for the players involved in the game, we show how these can be learned automatically from a body of data. We further show new definitions of quasi-polynomially computable approximate equilibria and show how to efficiently and approximately compute them. We apply these methods to the behavior of real world terrorist groups. Continuing with our work on adversarial models, we study methods to disclose information publicly in order to shape the adversary's actions to our advantage. We further develop methods to quantify and reduce lethality of networks. Our work develops and further studies the use of game-theoretic frameworks to the study of security games, online games, as well as games involving diverse terrorist groups including Lashkar-e-Taiba and the Indian Mujahideen.

I. Scientific Progress and Accomplishments

During this project, we made important contributions in the following 5 broad areas.

Data-Driven Learning of Utilities in Games. Given a body of data about the actions of one or more players in a given real-world scenario (e.g. terrorist attacks), can we learn the payoffs of different actions for a given player? In order to achieve this, we set the problem up as a combinatorial optimization problem in which constraints are generated under the assumption that the players prior actions (in the dataset) are based on what is best for them. The variables in the constraints correspond to payoffs for the players. We developed multiple ways of estimating the payoffs by looking at different types of solutions to the constraints and showed that our method to estimate payoffs performs well in predicting future actions by the actors, suggesting that the payoffs have been correctly learned.

Reasoning with Multiple Payoff Measures. Human experts' assessments of the values of utilities for different players in a game scenario can vary dramatically from one expert to another. In this work, we propose new notions of equilibria that build on Shapley's notion of approximate equilibria and develop quasi polynomial time approximation schemes to compute these equilibria

approximately, but efficiently. We applied these equilibria to the case of a real-world 5 player game involving the terrorist group Lashkar-e-Taiba.

Game Theoretic Adversarial Reasoning about Geospatial Phenomena. In past work we had developed the concept of geo-spatial abduction problems or GAPs and successfully used GAPs to predict locations of IED weapons caches in Baghdad and the locations of high value targets in Helmand and Kandahar provinces of Afghanistan. In this work, we were interested in the question: what if they adversary knew these IED cache (or HVT) location prediction methods and took evasive action. We developed a game theoretic model of the adversary's behavior and developed actions that the defender could take in order to counteract such adversarial behavior. We tested the theory on 21 months of data with very strong positive results.

Temporal Probabilistic Adversarial Reasoning. We developed the notion of Annotated Temporal Probabilistic (APT) logic to reason about adversaries such as terrorist groups. An APT-logic program consists of a set of rules of the form *If condition C is true at time t, then player P will take action A at time $(t + \delta)$ with a probability in the interval $[L, U]$.* Using APT logic programs, we can express rules about how a condition in a terror group's environment today will lead (potentially) to a future action carried out by the group. We applied this work to modeling the terrorist group *Indian Mujahideen*.

Reasoning about Adversarial Terrorist Networks. In order to destabilize a terrorist network, we must first be able to measure its lethality. We developed the first ever theoretical model of the lethality of a terrorist network and showed that this model can be used to accurately predict the number of attacks a particular network structure will carry out in the future. We then developed an algorithm to predict a probability distribution on who will replace a terrorist who is removed from a network (e.g. by a capture/kill operation) as well as a probability distribution over the space of network structures formed when a given set of terrorists are removed. Based on this we propose a novel algorithm to identify a set of k terrorists whose removal would minimize the operational efficiency of the network. A prototype system called STONE was built to achieve this purpose.

II. CONTRIBUTIONS

II.A Automated Data-Driven Learning of Payoffs

Virtually almost all work in game theory starts with a payoff matrix. In his pioneering study of conflict, Schelling [Schelling 1960] starts out with a payoff matrix for virtually every scenario. Unfortunately, getting a payoff matrix poses an enormous challenge in many real-world strategic games. In this paper, we answer the following question: Given a body of historical data about the interactions of multiple players, is there a way to learn a payoff matrix? In order to answer this question, we make certain assumptions:

Time Discounting. We believe that players are more likely to be influenced by “recent” history as opposed to events from a distant past. In order to model this, we developed a notion of time-discounted regret.

No Correlated Equilibria, Short Histories. We do not assume correlated equilibria nor do we assume the existence of a signaling mechanism. When long histories are available and some extra assumptions are made, game play can converge to correlated equilibrium even without a signaling mechanism. However, our real-world applications have short histories for which convergence cannot be assumed.

Bounded Rationality. Unlike past work on inverse reinforcement learning that assume fully rational agents, we assume bounded rationality, i.e. players take actions whose payoffs are within ϵ percent of the action with best response payoff.

No Knowledge of Outcomes. We assume that we only know the game history but nothing about outcomes.

Best Response. We assume that all players have complete knowledge of the history of past events and that in each time period, players choose an action that is an approximate best response to the history (subject to the bounded rationality assumption).

We define constraints whose variables represent the payoffs for each player under each joint action. Our constraints informally state that at each time point t in the past, each player i chose to perform the action for which he had the maximal expected time-discounted regret prior to time t . We also interpret these constraints as a myopic best response to the state of the world. This leads to a set of constraints with many possible solutions. We define three heuristics to estimate payoffs.

(1) **Centroid Solution (CBS).** In CBS, the (approximate) centroid of the constraint polytope is picked as the solution.

(2) **Soft Constraints Approach (SCA).** In SCA, we allow the rationality constraints to be violated but penalize such violations in the objective function used.

(3) **SVM-based Method (SVMM).** In SVMM, we propose a heuristic method to map the payoff inference problem onto a support vector machine and build a separator that captures the payoff function we wish to learn.

We implemented CBS, SCA, SVMM, as well as the recent ICEL algorithm (Inverse Correlation Equilibrium Learning) for comparison due to Waugh et al. in 2011. We compared all 4 algorithms w.r.t. solution quality and run-time. On synthetic data where we knew the ground truth (because we generated player behavior using known payoff functions), we showed that

SVMM outperforms both CBS and SCA w.r.t. both solution quality and run-time. We also compared CBS, SCA, and SVMM on two real-world data sets:

- (i) the Minorities at Risk Organizational Behavior (MAROB) dataset the contains data on terrorist group behaviors and related government actions and
- (ii) a much more fine-grained data set about the behavior of the terrorist group Lashkar-e-Taiba (LeT).

Again, SVMM outperformed CBS and SCA. We then ran experiments comparing SVMM with ICEL. When we compare the ability of SVMM with that of ICEL to predict true behaviors from learned payoffs on the MAROB data, SVMM’s ability to predict behavior from the learned payoffs was much better than that of ICEL (median Spearman Correlation Coefficient of 0.7 for SVMM, compared to just 0.114 for ICEL).

II.B Game-Theoretic Reasoning with Multiple Payoffs

The research reported here was motivated by a concrete application: how can countries trying to rein in the terrorist group Lashkar-e-Taiba (LeT) come up with policies against them, especially if these policies need to be coordinated? In the case of a five-player game that we formulated for LeT (presented later), there were wide variations of opinion among experts on what to do about LeT with respect to, for instance, whether India should carry out covert action, carry out coercive diplomacy, propose peace talks, or just keep the status quo. Likewise, the United States has historically had multiple opposing viewpoints on whether to continue financial (development and military aid) to Pakistan, whether to carry out covert action against LeT, or do nothing.

Analyzing the benefits of these actions even in the case of a single actor (e.g., only India or only the United States) has proven challenging. The main contribution of this work is a multiplayer, game-theoretic framework in which this specific problem can be solved.

However, we wanted to come up with a general solution—one that is applicable to many different settings. For instance, there are many applications where the “payoff matrices,” usually one of the very first things needed in any game-theoretic framework, cannot be specified with accuracy. When asked about payoffs, multiple experts might express substantial disagreements. This is what happened with our LeT application.

Approximate Equilibria for Multiplayer Games with Vector Payoffs: Games with multiple payoffs were introduced by Shapley in 1959. Shapley called them *vector-valued games*, and they have been extensively studied under various other names such as *multicriteria games* and *multiobjective games*. Unfortunately, for real-world applications such as the LeT application motivating this research, the computational cost of these past methods is too high. To address this, we introduce a novel combination of vector-valued games and approximate equilibria and

define new types of approximate equilibria for games with multiple players and multiple payoff matrices. We design algorithms for computing such equilibria for zero-sum games and games of low rank. For the case of rank-1 games, we give a structural result and use it to design a simple algorithm for such games. For general games, we give an extension of the approximation lemma of Althofer for simultaneous games with multiple payoff functions (SGMs) and use it to design a quasipolynomial time approximation scheme (QPTAS) when the number of players in a game is constant (which is the case for our LeT game).

Application of PREVE to Generate Policies to Reduce Terror Acts by LeT: We apply this theory to a real-world application in which there are five parties including four governmental entities and the terrorist group LeT. The goal was to understand whether there were any pure (or mixed) equilibria in which the group's terrorist acts could be significantly reduced. The five players considered are the United States, India, the Pakistani military, the Pakistani civilian government, and the terrorist group LeT.

When it comes to the application of game-theoretic reasoning to international strategic elements with both state and nonstate actors, the situation becomes much more complex because identifying the payoffs for different players is an enormous challenge and experts vary widely on what these payoffs are. To address this application, we asked three internationally acknowledged world experts to give us payoff matrices—and we received three payoff matrices with substantial differences between them. Leveraging the theoretical constructs and results described earlier, we built a prototype policy recommendation engine based on vector equilibria called PREVE and used it to identify approximate equilibria in the multiple payoff game induced by the three expert payoff matrices. The equilibria suggest that two actions will help cut down terror attacks by LeT: (i) discontinuing US aid to Pakistan, and (ii) covert action and coercive diplomacy by India toward Pakistan.

II.C Game-Theoretic Geospatial Abduction with Adversarial Models

Geospatial Abduction Problems (GAPs) were previously introduced by us to find a set of locations that “best explain” a given set of locations of observations. We call these inferred sets of locations “explanations.” There are many such applications in a wide variety of domains (military, criminology, conservation, epidemiology etc.) for which GAPs are useful. Our past work applied it to the prediction of locations of IED weapons caches in Baghdad and high value targets in Afghanistan.

Our past work assumed that the adversary (the “bad guy” or the entity that wishes to evade detection) does not reason about the agent (the “good guy” or the entity that wants to detect the adversary). Despite this significant omission, we were able to accurately predict the locations of weapons caches in real-world data about IED attacks in Baghdad.

In this work, we introduce adversarial geospatial abduction problems where both the agent and the adversary reason about each other. Specifically, our contributions are as follows.

- (1) We axiomatically define reward functions to be any functions that satisfy certain basic axioms about the similarity between an explanation chosen by the adversary (e.g., where the serial killer lives and works or where the insurgents put their IED caches) and define notions of expected detriment (to the adversary) and expected benefit (to the agent).
- (2) We formally define the Optimal Adversary Strategy (OAS) that minimizes chances of detection of the adversary’s chosen explanation and the Maximal Counter-Adversary strategy (MCA) that maximizes the probability that the agent will detect the adversary’s chosen explanation.
- (3) We provide a detailed set of results on the computational complexity of these problems, the counting complexity of these problems, and the possibility of approximation algorithms with approximation guarantees for both OAS and MCA.
- (4) We develop Mixed Integer Linear Programming algorithms (MILPs) for OAS and two algorithms, MCA-LS and MCA-GREEDY-MONO, to solve MCA with certain approximation guarantees. MCA-LS has no assumptions, while MCA-GREEDY-MONO assumes monotonicity.
- (5) We develop a prototype of our MILP algorithms to solve the OAS problem, using our techniques for variable reduction on top of an integer linear program solver. We demonstrate the ability to achieve near-optimal solutions as well as a correct reduction of variables by 99.6% using a real-world dataset.
- (6) We develop a prototype implementation that shows that both MCA-LS and MCA-GREEDY-MONO are highly accurate and have very reasonable time frames. Though MCA-GREEDY-MONO is slightly faster than MCA-LS, we found that on every single run, MCA-LS found the exact optimal benefit even though its theoretical lower-bound approximation ratio is only $1/3$. As MCA-LS does not require any additional assumptions and as its running time is only slightly slower than that of MCA-GREEDY-MONO, we believe this algorithm has a slight advantage.

II.D Adversarial Reasoning about Network Behavior

In 2008, A Revolutionary Armed Forces of Colombia, or FARC, commander named Nelly Avila Moreno (a.k.a. Karina) turned herself in to Colombian authorities in response to the announcement of a \$1 million award for her arrest. Unlike expensive and risky operations to capture terrorists (such as Al Qaeda’s Khalid Sheikh Mohammed and the Kurdish PKK’s Abdullah Ocalan), Karina had been captured at minimal expense in terms of both financial cost and lives put at risk. The Shaping Terrorist Organization Network Efficiency, or STONE, software platform we developed is designed to identify a set of key operatives in a terrorist network whose removal would maximally defang the organization through a variety of reward

programs and capture operations. STONE answers who should be the targets of the reward program and if a government wishes to destabilize a terrorist network and have funds to remove k people, which k people it should target.

Such removal operations are essential to international security. Though world governments spent more than \$70 billion fighting terrorism from 2001 to 2008, reducing the number of transnational attacks by 34%, there was a net increase in terrorism fatalities by 67 deaths per year during the same period.

Counterterror efforts have weighed strategic actions that try to address the root causes of terrorism by providing incentives to all parties to reach agreement, with many conducting strategic studies to reduce terrorism. However, strategic defeat of terrorist organizations is rare, despite some notable successes (such as the Provisional IRA in Ireland and Aum Shinrikyo in Japan).

As a consequence, tactical actions aimed at destabilizing terror networks are still necessary today. STONE uses three novel algorithms:

Terrorist Successor Problem (TSP). When a terrorist r is removed, it identifies the probability that r is replaced by another terrorist v ;

Multiple Terrorist Successor Problem (MTSP). When multiple terrorists are removed from a network, it identifies the new possible networks that might arise, together with an associated probability distribution; and

Terrorist Network Reshaping Problem (TNRP). It uses the results generated by MTSP to identify a set of k terrorists to remove from the network so as to minimize the expected efficacy of the resulting network. In the terrorist networks we consider, each vertex (such as Abdullah Azzam) could have many properties, including a status property specifying if he is alive, dead, or jailed; a role property specifying whether he is a fundraiser, ideologue, or recruiter. In addition to the vertex properties, we also consider hostility of a vertex toward the West, capability to launch attacks, and blowback if captured. A property labeling $p(v,p)$ tells us the value of property p for vertex v .

However, STONE can work with any set of properties, not just these. Each vertex in the network also has a rank—coded from 1 to 10, with 10 being the highest; multiple people may have the same rank. We developed the STONE algorithms and tested them with the help of military, policy, and terrorism experts.

The results were extremely promising. The STONE software was transitioned to ARL and the project was briefed at various security agencies in the Washington DC area.

II.E Game-Theoretic Disclosure Models

In our past work, we developed temporal-probabilistic rules that are predictive of attacks by terrorist groups such as Lashkar-e-Taiba and Indian Mujahideen. One question that was raised by Nobel Laureate Tom Schelling was the following: how does the disclosure of these rules affect counter-terrorism operations?

For this purpose, we developed a two player game-theoretic model in which the two players, naturally, are a defender (us) and an adversary (the terrorists). The defender has learned a set of behavioral rules about the adversary, but must decide which, if any, of them to disclose publicly.

Underlying the game are four factors about the adversary. Mounting an attack of a certain type has a *cost* to the adversary as well as a *benefit*. The cost can be understood as the resources the organization must apply to carry out an operation, which could include operatives killed, as well as money and time spent on planning, logistics, and surveillance. Likewise, each attack has some benefits to the attacker. The 2008 Mumbai attacks showed that the immense media coverage surrounding the attacks – more or less nonstop on all major TV and news channels worldwide for 3 days – was a huge benefit to LeT. When terrorists successfully target symbols of their hated enemy (such as the Pentagon and World Trade Center in the United States or the Taj Hotel and Red Fort in India) they show their supporters and sponsors that they are strong and capable, increasing their fund-raising and recruitment from the masses of radicalized individuals. Terrorist groups can derive strategic benefits from mounting certain attacks. As former FBI Director Louis Freeh says in the foreword to a recent book on the IM, the Iranian Revolutionary Guard Corp (IRGC) was responsible for the 1996 attack on US servicemen in the Khobar Towers in Saudi Arabia which strained U.S.-Saudi relations. Over a decade earlier the IRGC's close ally Hezbollah struck the U.S. Marine barracks in Beirut, causing the United States to withdraw from Lebanon and allow Hezbollah and its Syrian and Iranian sponsors to dominate that country. The third factor is *probability of success*. Terrorists are not stupid – they do not want to die (or commit resources) without being successful. If they must die, they prefer to do so with a bang that takes many of their enemies with them. The fourth factor is *defendability* – the defender can defend well against certain types of attacks (e.g. attacks against transportation sites like airports and train stations) and less well against other types of attacks (e.g. outdoor public markets for example).

Using these four factors as input, we first put ourselves in the shoes of the attacker. We asked ourselves: suppose a set D of behavioral rules (known by the defender about the attacker) are also known to the attacker, how would the attacker then act? What would he attack and what would he not attack? We assumed that if the defender reveals the set D of rules, then he would also take actions on the ground to convince the attacker that carrying out those types of attacks would have a much lower probability of success. For instance, there would be no point if the TSA asserted that they knew that Al-Qaeda terrorists were using shoe bombs, but did not check shoes systematically. In this case, the TSA's knowledge of the rule would not serve as a

deterrent. For a disclosure to serve as a deterrent, concrete action on the ground is needed to convince the adversary that we are serious. Based on this, our mathematical model computed the attacks the adversary would carry out.

Once we understand what the adversary might do when different sets $\mathbf{D}$ of rules are disclosed, the defender can identify the subset $\mathbf{D}_{\min}$ of the true set $\mathbf{D}_{\text{true}}$ of learned rules that he should disclose publicly that minimize the expected impact of the attacker's attacks.

We were able to mathematically show the following significant results:

- *Disclosing no rules about an adversary is not always the best strategy.*
- *Disclosing everything we know about an adversary is not always the best strategy.*

In general, our mathematical results to date show conclusively that the disclosure of behavioral models of terrorist groups must be strategic. We must carefully “game out” the benefits of disclosure versus the risks, and identify the set of disclosures that best deter the attacker. Of course, these results are mathematical. Most researchers, including us, have no way to explicitly engage the adversary. We do not know if they are listening carefully to our disclosures and if their actions (or non-actions) are due to our disclosures or due to something else, although research on how terrorist groups acquire information and learn suggests that effective terrorist groups systematically study the activities and capabilities of the states in which they operate and target. It is understandable that intelligence agencies believe that keeping their findings secret will surprise the adversary. But our results indicate that this may not always be the optimal policy. Strategic disclosures have played a strong deterrent role for centuries, and their value in the fight against terrorism has been previously noted. New technology such as ours will provide better tools to assess the consequences of such strategic disclosures so as to make them a potent weapon in today's battle against terrorism.

II.F Other Contributions

In addition, we made a number of other important contributions in this project which are briefly summarized below.

- We considered a family of repeated bilateral games of incomplete information called “choice selection processes”, in which players may share certain goals, but are essentially self-interested. We focused on road selection problems and described several possible models of human behavior that were inspired by behavioral economic theories of people's play in repeated interactions. The results revealed that combining a hyperbolic discounting model of human behavior with a social utility function yields the best results. The prediction rate of all methods is summarized in the following table:

Method	Prediction
Rational	45%
Exponential Smoothing	62.86%
Hyperbolic Discounting	64.17%
Short Memory	60.33%

- We modeled people's decision-making strategies during negotiation in two separate works. The negotiation is conducted in Colored Trail games. The Colored Trail (CT) family of games was designed by Kraus and Grosz in order to inhabit the middle ground between complex, real world task domains and the payoff-matrix abstractions of behavioral economics games. It thus provides a realistic but modeling-tractable setting.
 - In a first work we studied a setting which involved an alternating-offer protocol that allowed parties to choose the extent to which they honored each of their agreements during the negotiation. The agent tried to predict how people reciprocate their actions over time despite the scarcity of prior data of their behavior across different cultures. Our methodology addresses this challenge by using classical machine learning techniques to predict how people respond to offers and the extent to which they fulfill agreements. We built different models for each culture. We used features such as past reliability and agreement benefit. The accuracy is as follows:

Model / Population	Israel	Lebanon	U.S.
Reliability	84.127 (PUK)	97.5 (PUK)	85.7143 (PUK)
Acceptance	70.3883 (KNN:k=19)	75(J48)	73.5016 (KNN; k=9)

In a second work we considered situations of repeated negotiation in incomplete information settings where players need to decide strategically whether to reveal information during the negotiation process. We used classical machine learning techniques to predict how people make

and respond to offers during the negotiation, how they reveal information and their response to potential revelation actions by the agent. We used features such as (i) the number of chips needed to get from the current position of the player to the goal; (ii) the difference between the chips in the possession of a player and the chips it is given at the onset of the game. The accuracy of the prediction is as follows:

Model	Predictor	Accuracy
Accepting proposals	SVM (linear kernel)	71%
Proposals	multi-class logistic regression	68%
Revelations	multi-class logistic regression	72%

- We consider the problem of designing automated strategies for interactions with human subjects, where the humans must be rewarded for performing certain tasks of interest [Azaria et al 2012b]. We focus on settings where there is a single task that must be performed many times by different humans (e.g. answering a questionnaire), and the humans require a fee in order to perform the task. We wanted to predict the price that needed to be offered to a worker so that he would agree to work on the task. We tested two methods for price elicitation: (1) Vickrey auction: A Vickrey auction is a sealed bid auction where each worker submits a bid and the worker with the lowest bid performs the task and is paid the amount requested by the second lowest bid. (2) We assume that the portion of workers that would accept an offer of x follows a sigmoidal distribution (in x). We approximated the distribution by choosing several points, then sampled a subset of workers and obtained their acceptance fraction for these points (different workers for different points) and interpolated the sigmoid from these values. We tested the prediction of both methods. The Mean Squared Error (MSE) of the Vickrey based method was 0:46, which is extremely high, while the Sigmoid based method's error was only 0:18. This result suggests that price elicitation by sampling a subset of workers and interpolating a sigmoid is much more accurate than using a Vickrey Auction for the same purpose.
- An important way to learn new actions and behaviors is by observing others, and several evolutionary games have been developed to investigate what learning strategies work best and how they might have evolved.
 - We have developed an extensive set of mathematical and simulation results for Cultaptation, which is one of the best-known such games.

- Our results include a formula for measuring a strategy's expected reproductive success, algorithms to compute near-best-response strategies and near-Nash equilibria, and techniques for efficient implementation of those algorithms.
- Our experimental studies provide strong evidence for the following hypotheses:
 - The best strategies for Cultaptation and similar games are likely to be conditional ones in which the choice of action at each round is conditioned on the agent's accumulated experience. Such strategies (or close approximations of them) can be computed by doing a lookahead search that predicts how each possible choice of action at the current round is likely to affect future performance.
 - Such strategies are likely to exploit most of the time, but will have ways of quickly detecting structural shocks, so that they can switch quickly to innovation in order to learn how to respond to such shocks. This conflicts with the conventional wisdom that successful social-learning strategies are characterized by a high frequency of innovation; and agrees with recent experiments by others on human subjects that also challenge the conventional wisdom.
- Until recently, game-tree pathology (in which a deeper game-tree search results in worse play) has been thought to be quite rare. We have performed an analysis that shows that every game should have some sections that are locally pathological, assuming that both players can potentially win the game. We also have developed a way to overcome game-tree pathology in some cases, by modifying the minimax algorithm to recognize local pathologies in arbitrary games, and cut off search accordingly (shallower search is more effective than deeper search when local pathologies occur). We have shown experimentally that our modified search procedure avoids local pathologies and consequently provides improved performance, in terms of decision accuracy, when compared with the ordinary minimax algorithm.
- We proposed an efficient agent for competing in Cliff-Edge (CE) and simultaneous Cliff-Edge (SCE) situations. In CE interactions, which include common interactions such as sealed-bid auctions, dynamic pricing and the ultimatum game (UG), the probability of success decreases monotonically as the reward for success increases. This trade-off exists also in SCE interactions, which include simultaneous auctions and various multi-player ultimatum games, where the agent has to decide about more than one offer or bid simultaneously. Our agent competes repeatedly in one-shot interactions, each time against different human opponents. The agent learns the general pattern of the population's behavior, and its performance is evaluated based on all of the interactions in which it participates. We propose a generic approach which may help the agent compete against unknown opponents in different environments where CE and SCE interactions exist, where the agent has a relatively large number of alternatives and where its achievements in the first several dozen

interactions are important. The underlying mechanism we propose for CE interactions is a new meta-algorithm, deviated virtual learning (DVL), which extends existing methods to efficiently cope with environments comprising a large number of alternative decisions at each decision point. Another competitive approach is the Bayesian approach, which learns the opponents' statistical distribution, given prior knowledge about the type of distribution. For the SCE, we propose the simultaneous deviated virtual reinforcement learning algorithm (SDVRL), the segmentation meta-algorithm as a method for extending different basic algorithms, and a heuristic called fixed success probabilities (FSP). Experiments comparing the performance of the proposed algorithms with algorithms taken from the literature, as well as other intuitive meta-algorithms, reveal superiority of the proposed algorithms in average payoff and stability as well as in accuracy in converging to the optimal action, both in CE and SCE problems.

- Corruption frequently occurs in many aspects of multi-party interaction between private agencies and government employees. Past works studying corruption in a lab context have explicitly included covert or illegal activities in participants' strategy space or have relied on surveys like the Corruption Perception Index (CPI). This work studies corruption in ecologically realistic settings in which corruption is not suggested to the players a priori but evolves during repeated interaction. We ran studies involving hundreds of subjects in three countries: China, Israel, and the United States. Subjects interacted using a four-player board game in which three bidders compete to win contracts by submitting bids in repeated auctions, and a single auctioneer determines the winner of each auction. The winning bid was paid to an external "government" entity, and was not distributed among the players. The game logs were analyzed posthoc for cases in which the auctioneer was bribed to choose a bidder who did not submit the highest bid. We found that although China exhibited the highest corruption level of the three countries, there were surprisingly more cases of corruption in the U.S. than in Israel, despite the higher PCI in Israel as compared to the U.S. We also found that bribes in the U.S. were at times excessively high, resulting in bribing players not being able to complete their winning contracts. We were able to predict the occurrence of corruption in the game using machine learning. The significance of this work is in providing a novel paradigm for investigating covert activities in the lab without priming subjects, and it represents a first step in the design of intelligent agents

III. STUDENTS/POSTDOCS INVOLVED

Students: M. Purohit, A. Sawant, B. Schuetzle, P. Shakarian, B. Wilson, R. Carr, E. Raboin, A. Sliva, P. Roos

Postdocs: A. Mannes, E. Serra, F. Spezzano, A. Zukerman, C. Molinaro, P. Roos

PAPERS PUBLISHED/ACCEPTED/SUBMITTED

1. G. Simari, J. Dickerson, A. Sliva and V.S. Subrahmanian. "Parallel Abductive Query Answering in Probabilistic Logic Programs", *ACM Transactions on Computational Logic*, 14(2):12, 2013.
2. Puzis, Rami, Manish Purohit, and V. S. Subrahmanian. "Betweenness computation in the single graph representation of hypergraphs." *Social Networks* 35.4 (2013): 561-572.
3. P. Shakarian and V.S. Subrahmanian. Region-based Geospatial Abduction with Counter-IED Applications, in: Counterterrorism and Open Source Intelligence (ed. U.K. Wiil), Springer Verlag Lecture Notes on Social Networks, pp. 105-142.
4. Brandon Wilson, Inon Zuckerman, Austin Parker, and Dana Nau. Improving Local Decisions in Adversarial Search. Proc. European Conference on Artificial Intelligence (ECAI), 2012.
5. Ryan Carr, Eric Raboin, Austin Parker, and Dana Nau. Theoretical and Experimental Analysis of an Evolutionary Social-Learning Game. Submitted for journal publication, 2012.
6. Azaria A, Aumann Y, Kraus S. Automated Strategies for Determining Rewards for Human Work, AAAI 2012.
7. Galit Haim, Yakov Gal, Sarit Kraus, A Cultural Sensitive Agent for Human-Computer Negotiation. *Proc. of AAMAS 2012*.
8. N. Peled, Y. Gal and S. Kraus. Learning to Reveal Information in Repeated Human-Computer Negotiation, *Human-Agent Interaction Design and Models (HAIDM) workshop*, 2012.
9. M. Bitan, Y. Gal, E. Dokow and S. Kraus. Social Rankings in Human-Computer Committees, *Human-Agent Interaction Design and Models (HAIDM) workshop*, 2012.
10. Asaf Frieder, Raz Lin, Sarit Kraus. Agent-human Coordination with Communication Costs under Uncertainty (Extended Abstract), *Proc. Of AAMAS 2012*.
11. Azaria A, Rabinovich Z, Kraus S, Goldman CV, Gal Y. Strategic Advice Provision in Repeated Human-Agent Interactions, AAAI 2012.
12. F. Parisi, A. Sliva and V. S. Subrahmanian. A Temporal Database Forecasting Algebra, *International Journal of Approximate Reasoning*, Vol. 54, Nr. 7, pages 827-860, 2013.
13. C. Molinaro, A. Sliva, and V.S. Subrahmanian. Super-solutions: Succinctly Representing Solutions in Abductive Annotated Probabilistic Temporal Logic, submitted to *ACM Transactions on Computational Logic* in March 2012, revised July 2013 in accordance with reviewer comments.
14. A. Sawant, J. Dickerson, M.T. Hajiaghayi, and V.S. Subrahmanian. PREVE: A Policy Recommendation Engine based on Vector Equilibria, *Proc. 2013 Intl. Conf. on Foundations of Open-Source Intelligence and Security Informatics*, Aug 2013.
15. Eric Raboin, Petr Svec, Dana Nau, and Satyandra K. Gupta. Model-Predictive Target Defense by Team of Unmanned Surface Vehicles Operating in Uncertain Environments. *International Conference on Robotics and Automation (ICRA)*, 2013.

16. R. Carr, E. Raboin, and D. Nau. Theoretical Analysis and Evolutionary Simulation of Social-Learning Strategies. *Journal of Artificial Societies and Social Simulation (JASSS)*, under review.
17. Patrick Roos. *Evolutionary Game Theoretic Modeling of Decision Making and Culture*. Ph.D. Dissertation, 2013.
18. N. Peled, M. Bitan, Y. Keshet and S. Kraus. Predicting Human Strategic Decisions Using Facial Expressions, *Proc. Of IJCAI 2013*.
19. A. Azaria, A. Richardson and S. Kraus. Autonomous Agent for Deception Detection, in *Proc. of HAIDM (In AAMAS) 2013*.
20. A. Rosenfeld, Z. Bareket, C. V. Goldman, S. Kraus, D. J. LeBlanc and O. Tsimhoni. Towards Adapting Cars to their Drivers, *AI Magazine*, 33(4):46-58, 2012.
21. Thanh H. Nguyen, Rong Yang, Amos Azaria, Sarit Kraus and Milind Tambe. Analyzing the Effectiveness of Adversary Modeling in Security Games, *Proc. Of AAAI 2013*.
22. A. Sawant, J. Dickerson, M. Hajiayaghi and V.S. Subrahmanian. Automated Generation of Counter-Terrorism Policies using Multi-Expert Input, *ACM Transactions on Intelligent Systems & Technology*, 6(6):44.
23. E. Serra and V.S. Subrahmanian. Should Behavioral Models of Terrorist Groups be Disclosed?, *IEEE Intelligent Systems*, Vol. 29, Nr. 4, 2014.
24. E. Serra and V.S. Subrahmanian. A Survey of Quantitative Models of Terror Group Behavior and an Analysis of Strategic Disclosure of Behavior Models, *IEEE Transactions on Computational Social Science* in March/April 2014.
25. F. Spezzano, V.S. Subrahmanian, and A. Mannes. Reshaping Terrorist Networks, *Communications of the ACM*, Vol. 57 No. 8, pp. 60-69, August 2014.
26. P. Shakarian, J. Dickerson, and V.S. Subrahmanian. Adversarial Geospatial Abduction Problems, *ACM Transactions on Intelligent Systems Technology*, Vol. 3, Nr. 2, Article Nr. 34, 35 pages, Feb. 2012.
27. A. Sawant, M. Hajiaghayi, S. Kraus, E. Serra, and V.S. Subrahmanian. PIE: A Data-Driven Payoff Inference Engine for Multi-Player Games with Short Histories and Bounded Rationality, submitted to a technical journal in Feb 2016.
28. Azaria, A. Richardson and S. Kraus. An Agent for Deception Detection in Discussion Based Environments. The 18th ACM conference on Computer-Supported Cooperative Work and Social Computing, 2015. <http://u.cs.biu.ac.il/~sarit/data/articles/pirate.pdf>
29. A. Azaria, Y. Aumann and S. Kraus. Automated Agents for Reward Determination for Human Work in Crowdsourcing Applications, *Journal of Autonomous Agents and Multi-Agent System*. 28(6): 934-955, 2014.
30. R. Azoulay, R. Katz and S. Kraus. Efficient Bidding Strategies for Cliff-Edge Problems, *Journal of Autonomous Agents and Multi-Agent Systems (JAAMAS)*, 28:290–336, (2014).
31. Y. Gal, A. Rosenfeld, S. Kraus, M. Gelfand, B. An and J. Lin. New Paradigm for the Study of Corruption in Different Cultures, *Proc. Of 2014 International Conference on Social Computing, Behavioral-Cultural Modeling, & Prediction (SBP14)*, 2014. <http://u.cs.biu.ac.il/~sarit/data/articles/corruption.pdf>

32. A. Rosenfeld, I. Zuckerman, E. Segal-Halevi, O. Drein and S. Kraus. NegoChat: A Chat-Based Negotiation Agent. AAMAS 2014.
33. Azaria A. and Kraus S. Advice Provision in Multiple Prospect Selection Problems, HAIDM (In AAMAS) 2013 (also appears as an extended abstract in AAAI 2013).
34. M. Bitan, Y. Gal, S. Kraus, E. Dokow and A. Azaria. Social Rankings in Human-Computer Committees, Proc. of AAAI 2013.
35. A. Rosenfeld, I. Zuckerman, A. Azaria, and S. Kraus. Combining Psychological Models with Machine Learning to Better Predict People's Decisions, Synthese, 189:81-93, 2012.