# PIE: A Data-Driven Payoff Inference Engine for Strategic Security Applications

Haipeng Chen, Mohammad T. Hajiaghayi, *Fellow, IEEE*, Sarit Kraus, Anshul Sawant, Edoardo Serra, V. S. Subrahmanian, and Yanhai Xiong

***Abstract*— Although most game theory models assume that payoff matrices are provided as input, getting payoff matrices in strategic games (e.g., corporate negotiations and counter-terrorism operations) has proven difficult. To tackle this challenge, we propose a payoff inference engine (PIE) that finds payoffs assuming that players in a game follow a myopic best response or a regret minimization heuristic. This assumption yields a set of constraints (possibly nonlinear) on the payoffs with a multiplicity of solutions. PIE finds payoffs by considering solutions of these constraints and their variants via three heuristics. First, we approximately compute a centroid of the resulting polytope of the constraints. Second, we use a soft constraint approach that allows violation of constraints by penalizing violations in the objective function. Third, we develop a novel approach to payoff inference based on support vector machines (SVMs). Unlike past work on payoff inference, PIE has the following advantages. PIE supports reasoning about multiplayer games, not just one or two players, it can use short histories, not long ones which may not be available in many real-world situations, it does not require all players to be fully rational, and it is one to two orders of magnitude more scalable than past work. We run experiments on a synthetic data set where we generate payoff functions for the players and see how well our algorithms can learn them, a real-world coarse-grained counter-terrorism data set about a set of different terrorist groups, and a real-world fine-grained data set about a specific terrorist group. As the ground truth about payoffs for the terrorist groups cannot be tested directly, we test PIE by using the payoffs to make predictions about the actions of the groups and corresponding governments (even though this is not the purpose of this article). We show that compared with recent work on payoff inference, PIE has both higher accuracy and much shorter runtime.**

***Index Terms*— Counter terrorism, game theory, payoff inference.**

## I. Introduction

GAME theory is a classic and powerful tool for modeling strategic behavior of a system of multiple agents who interact with each other. Almost all works in game theory start with a payoff matrix. In his pioneering study of conflict, Schelling [1] starts out with a payoff matrix for virtually every scenario. While getting a payoff matrix for a game is critical to modeling and understanding the behaviors of the involved agents, it unfortunately poses an enormous challenge in many real-world strategic games where such knowledge does not exist.

In this article, we address the payoff inference problem in a game of interacting players with our counter-terrorism application. Because counter terrorism is adversarial, we target noncooperative scenarios in order to answer the following question: *Given a body of historical data about the interactions of a set of noncooperative players, is there a way to learn a payoff matrix?* The approach in this article is partly motivated by our ongoing counter-terrorism research involving the terrorist group Lashkar-e-Taiba (LeT) (responsible for the 2008 Mumbai attacks)—given the history of interactions of the governments of Pakistan and India and LeT, we wish to understand their payoff [2], [3]. To answer these questions, we need a realistic game theoretic model that is able to characterize real-world game scenarios (e.g., in counter terrorism) while being robust to potential deviations from the game models. To this end, we consider the following properties in our model.

1) *Best Response:* Given the history of past events (i.e., choices of actions), in each time period, players choose an action that is an approximate best response to the history (subject to the bounded rationality property as described next).
2) *Bounded Rationality:* In real-world games, such as our counter-terrorism situations, decision makers are not likely to be fully rational but boundedly rational, i.e., players take actions whose payoffs are within $\epsilon$ percent of the action with best response payoff. Therefore, we assume bounded rationality, not full rationality.
3) *Time Discounting:* Intuitively, 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.
4) *No Correlated Equilibria, Short Histories:* A correlated equilibrium is a status where no player wants to deviate from the recommended strategy from a public signal (assuming the others do not deviate). When long histories are available and some extra assumptions are made [4], game play can converge to a correlated equilibrium even without a signaling mechanism. However,

Manuscript received April 10, 2019; revised November 3, 2019; accepted November 17, 2019. Date of publication January 8, 2020; date of current version February 24, 2020. This work was supported in part by Guggenheim Fellowship, NSF under Grant CCF:SPX 1822738 and Grant IIS:BIGDATA 1546108, in part by DARPA under Grant SI3CMD, in part by the UMD Year of Data Science Program Grant, and in part by the Northrop Grumman Faculty Award. *(Corresponding author: V. S. Subrahmanian.)*

H. Chen, V. S. Subrahmanian, and Y. Xiong are with Dartmouth College, Hanover, NH 03755 USA (e-mail: haipeng.chen@dartmouth.edu; vs@dartmouth.edu; yanhai.xiong@dartmouth.edu).

M. T. Hajiaghayi is with the Department of Computer Science, University of Maryland, College Park, MD 20742 USA (e-mail: hajiagha@cs.umd.edu).

S. Kraus is with the Department of Computer Science, Bar-Ilan University, Ramat Gan 5290002, Israel (e-mail: sarit@cs.biu.ac.il).

A. Sawant was with the Department of Computer Science, University of Maryland, College Park, MD 20742 USA. He is now with Google Inc., Mountain View, CA 94043 USA (e-mail: anshul.sawant@gmail.com).

E. Serra is with the Department of Computer Science, Boise State University, Boise, ID 83725 USA (e-mail: edoardoserra@boisestate.edu).

Digital Object Identifier 10.1109/TCSS.2019.2957178

many real-world applications have short histories for which convergence cannot be assumed. As a result, we do not assume correlated equilibria [4] or the existence of a signaling mechanism [5].

There has been extensive prior work on applying game-theoretic models to counter terrorism and more generally, security problems [3], [6]–[13]. These works usually assume that the payoff matrices of the players are known *a priori*, which is not the case in our counter-terrorism domain. The payoff inference problem has been studied in economics for various markets [14]–[17]. However, their focus is on modeling a particular market and then to use the domain-specific model for model fitting and regression. Therefore, their methods cannot be applied to our problem. Though there are several existing works on inverse reinforcement learning (IRL) which study the payoff inference problem, we will show the following.

1) Most of these works focus on single agent [18]–[20].
2) Some of these works focus on multiple agents in a cooperative setting [21], [22].
3) The works on multiagent noncooperative settings [23]–[26] generally [including those in 1) and 2)] target payoff inference problems defined on a Markov decision process (MDP), which is an overcomplication of our problem.

For payoff inference on games with multiple players, we show that past works lack at least one of these properties. One work that is closely related to our problem is [27]. However, we will show that the model formulated in [27] is computationally inefficient, as solving the model involves convex optimization. Moreover, another distinction of our model from existing works is that we take into consideration all the properties mentioned above (i.e., best response, time discounting, and no correlated equilibria), but existing works lack at least one of the properties in their game theory-based payoff inference model.

We provide two different formulations to model the payoff inference problem. The first formulation is built on top of the concept of regret in repeated decision making problems, where we define a set of constraints whose variables represent the tabular representation of 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 $t$. In the second model, we interpret these constraints as a myopic best response to the state of the world (i.e., a history of actions for all the players) with a (possibly) nonlinear function form representation of the payoff function. The two models lead to a set of constraints defined on the payoff functions.

To solve the above problems, we propose three approaches: centroid-based solution (CBS) and soft constraints approach (SCA) are devised for the first model, while SVM-based method (SVMM) is designed for the second model.

1) *CBS:* In CBS, the (approximate) centroid of the constraint polytope is picked as the solution.
2) *SCA:* In the SCA, we allow the rationality constraints to be violated but penalize such violations in the objective function.
3) *SVMM:* In the SVMM, we propose a heuristic method to map the payoff inference problem onto a support vector machine (SVM) [28] and build a separator that captures the payoff function we wish to learn.

We implemented CBS, SCA, and SVMM, as well as the inverse correlation equilibrium learning (ICEL) algorithm [27] on both synthetic data and two real-world data sets. We compared all four algorithms with respect to 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 the SVMM outperforms both CBS and SCA with respect to both solution quality and run time. We also compared CBS, SCA, and SVMM on two real-world data sets: 1) the Minorities at Risk Organizational Behavior (MAROB) data set [29] that contains data on terrorist group behaviors and related government actions and 2) a much more fine-grained data set [2] about the behavior of the terrorist group LeT.[1] Again, the SVMM outperformed CBS and SCA. We then ran experiments comparing the SVMM with ICEL. When we compare the ability of the 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 ICEL (median Spearman correlation coefficient of 0.7 for the SVMMcompared with just 0.114 for ICEL).

## II. Related Work

A major driver for this article is counter-terrorism applications. The development of game-theoretic methods to analyze terrorist behavior and organization has been pioneered in [6] and [7] and subsequently adopted by others [8]–[13]. However, as described earlier, these approaches usually assume that the payoff matrices are known in advance by the decision makers, which might not be the case in many real-world problems such as the counter-terrorism applications motivating this article.

Economists have studied payoff inference problems for various markets [14]–[17]. However, their focus is on modeling a particular market and then to use various model fitting and regression methods to learn the best parameters. For instance, [16] studies the effect of land use regulations on the midscale hotel market. In our problem, the decision making process is in an interactive environment [30], [31] with multiple agents (i.e., governments and terrorist groups) as opposed to the single agent scenario in these works. Therefore, these lines of research cannot be directly applied to our problem.

[1]As no ground truth exists about payoffs for real-world players in the MAROB and LeT data sets, we learned player payoffs from a training data set and then validated them on a separate validation data set by making predictions based on learned payoffs. We emphasize the fact that this article is not about prediction, but about learning payoffs in order to understand group behavior. The goal is to understand the payoff structure for different players for different strategies, so diplomats and counter-terrorism agencies can shape policies toward the terrorist groups. We use predictions solely to validate learned payoffs.

IRL [18] learns payoffs of a single agent operating in a given (usually Markovian) environment. Reference [18] addresses the problem of learning a reward function by observing the behavior of MDPs. However, they and a series of subsequent works [19], [20] assume a single rational agent in a given environment. Some recent works have focused on multiagent IRL (MIRL). In contrast to the scenario we consider in which players are noncooperative, [21] and [22] study the problem of learning player payoffs in the presence of a centralized coordinator. While several other works [23]–[26] study the noncooperative setting, a major difference between the problems addressed in these works is that they focus on noncooperative games defined on a MDP, while the problem targeted in this article is a one-shot decision making problem (i.e., the decision at the current time step does not affect the decisions at future time steps). Therefore, the approaches to MIRL are distinct from this article. Perhaps, the prior work that is closest to our problem setting is by Waugh *et al.* [27] who proposed an approach to predict player behavior when no payoff matrix is available. A convex optimization formulation finds a maximum entropy solution to find the predicted distribution over joint actions. Finally, payoffs are computed by using the dual of the above optimization convex problem. However, this approach suffers from the computational complexity brought by the convex program formulation.

The assumption of equilibrium is common to most work on MIRL and other payoff learning methods. However, decision theory strongly suggests that human players do not follow equilibrium strategies, even when the equilibrium is unique (which is also rare in real-world problems) [1]. However, decision theory does highlight the importance of recency [32] and regret in human decision making. Anticipated regret is considered an important determinant of choice behavior [33]–[35]. These aspects form the basis for PIE's time-discounted regret minimization and myopic best response with exponentially decaying state. Thus, PIE differs from existing work on payoff inference in that we assume myopic rationality and not global rationality (equilibrium). We also assume simple game play dynamics inspired by relevant work from decision theory. In addition, we develop a fast and practical data analytic approach compared with the more theoretical approach taken by most machine learning articles. We show this with experiments on two real-world data sets and show superior performance compared to Waugh *et al.* [27] who only study a very small, toy example.

## III. Payoff Inference Model

In this section, we first introduce the preliminaries of the game model, followed by two different formulations of the payoff inference models. The first model is based on the idea of "regret minimization" with a tabular representation of the payoff matrices. The second model, which is motivated by our counter-terrorism application on LeT, represents the payoff matrices in a function approximation form with respect to a "time-weighted history."

### A. Preliminaries

Let $[N] = \{1, \dots, N\}$ be a set of players. We assume that each player $i$ has an associated set $A_i$ of actions that it can take. Let $\mathcal{A} = A_1 \times \cdots \times A_N$ denote the set of all possible joint actions. Given a joint action $a \in \mathcal{A}$, $a_i$ is the action of player $i$ and $a_{-i}$ is the joint action of all other players. Let $u_i$ be an unknown payoff function: $u_i : \mathcal{A} \to [0, 1]$. $u_i(a)$ is the payoff of joint action $a$ for player $i$. Let $\mathcal{U} = \{u_1, \dots, u_N\}$ be the set of all (as yet unknown) payoff functions, where $u_i$ is the payoff function for player $i$. Let $[T] = \{1, \dots, T\}$ be a set of past time points.

Let $m = \sum_{i\in[n]} |A_i|$ be the total number of actions for all players in the game. We encode a joint action as an $m$-dimensional binary vector. Each action for a player $i$ is indexed from $\sum_{j\in\{1\dots,i-1\}} |A_j| + 1$ to $\sum_{j\in\{1\dots i\}} |A_j|$ in a fixed but arbitrary order. In other words, the first $|A_1|$ entries in the vector describe the actions for the first player, the next $|A_2|$ entries describe the actions for the second player, and so forth. Let $v$ be an encoding for $a \in \mathcal{A}$. If, in the joint action represented by $a$, player $i$ plays action $a_i \in A_i$ at time $t$, then, and only then is $v[a_i] = 1$, otherwise $v[a_i] = 0$.

*Example 1:* Suppose we have two players 1 and 2 and suppose $A_1 = \{a, b, c\}$ and $A_2 = \{a, e\}$ are the actions they can perform. Then, the dimensionality of a joint action is 5 and an example of the vector representation of a joint action is

| pl-1 | pl-1 | pl-2 | pl-2 | pl-2 |
|---|---|---|---|---|
| a | b | c | a | e |
| 1 | 0 | 0 | 0 | 1 |

The first row is the player's ID and the second row is the action name. Here, the 5-D vector (1,0,0,0,1) tells us that in this joint action, player 1 performed action $a$ and player 2 performed action $e$.

A *history* is a sequence $H^\tau = \langle a^1, \dots, a^\tau \rangle$, where $a^t$ is the vector of joint actions taken at time $t \in [T]$. We represent the history of a game as a matrix, $H$, where $H[t, a] = 1$ if and only if player $i$ plays action $a \in A_i$ at time $t$. Thus, $H_t$, the $t$th row of the history matrix represents the joint action taken by all players at time $t$. Likewise, the $i$'th column of $H$ tells us that actions taken by player $i$ at each time point.

*Example 2:* Suppose we have two players $[N] = \{1, 2\}$; player 1 is a terror group and player 2 is the government. Assume that the players' actions are pe_g ("political engagement with the government") for player 1 and pe_tg ("political engagement with the terror group") for player 2. Each of these variables has three possible levels of intensity (low, medium, and high). Therefore, player 1 has three actions, pe_g($l$), pe_g($m$), and pe_g($h$), corresponding to the three levels of intensity of this action, and similarly, player 2 has three actions pe_tg($l$), pe_tg($m$), and pe_tg($h$). Let indices of the actions pe_g($l$), pe_g($m$), and pe_g($h$) be 1–3, respectively, for player 1, and 4–6, respectively, for player 2. Suppose we have three years ($[T] = \{1, 2, 3\}$) of history in the following.

| $[T]$ | year | player 1 | player 2 |
|---|---|---|---|
| 1 | 2010 | pe_g($l$) | pe_tg($m$) |
| 2 | 2011 | pe_g($m$) | pe_tg($l$) |
| 3 | 2012 | pe_g($h$) | pe_tg($m$) |

Then, the history matrix $H$ is given by

$$\begin{pmatrix} 1 & 0 & 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 1 & 0 \end{pmatrix}.$$

*B. Regret-Based Payoff Inference*

In this section, we define the concept of time-discounted regret. Classical regret is defined with respect to a class $\Phi$ of modification functions. Each modification function $f \in \Phi$ is a mapping $f : A \rightarrow A$. Intuitively, a modification function suggests an alternative choice $f(a)$ for an action $a$. Instead of taking action $a$, the player takes action $f(a)$. As there are many ways in which a player could modify his choice, we consider a set $\Phi$ of modification functions. In the context of our running counter-terrorism example, the different modification functions might correspond to all feasible actions that could replace a given action $a$. The *regret* for a player $i$ is defined as

$$R_{i,\Phi}(t) = \max_{f \in \Phi} \sum_{\hat{t}=1}^{t-1} u_i\big(f\big(a_i^{\hat{t}}\big), a_{-i}^{\hat{t}}\big) - u_i(a^{\hat{t}})$$

where $u_i(f(a_i^{\hat{t}}), a_{-i}^{\hat{t}}) - u_i(a^{\hat{t}})$ is the difference in utility for player $i$, had he elected to take action $f(a_i^{\hat{t}})$ instead of whatever action he took at time $t$ in the past. The summation $\sum_{\hat{t}=1}^{t-1} u_i(f(a_i^{\hat{t}}), a_{-i}^{\hat{t}}) - u_i(a^{\hat{t}})$ reflects the total regret that player $i$ had with respect to his past actions, had he chosen to use modification function $f$ instead of whatever method he used to select his past actions. Had player $i$ used the modification function $f \in \Phi$ that maximizes this summation, then he would have gotten the maximal possible benefit, and the fact that he (maybe) did not use it what leads to this regret.

When determining what action to take, players in the real world are often more influenced by recent actions than by actions in the distant past. Our notion of *time-discounted regret* takes this into account by allowing a player to discount the past at a rate $\alpha$ subject to $0 < \alpha \leq 1$. After each time point, the "importance" of a past event is reduced by a factor of $\alpha$. The time-discounted regret is defined as follows:

$$TDR_{i,\Phi}(t) = \max_{f \in \Phi} \frac{\sum_{\hat{t}=1}^{t-1} \alpha^{t-1-\hat{t}}\big(u_i\big(f\big(a_i^{\hat{t}}\big), a_{-i}^{\hat{t}}\big) - u_i(a^{\hat{t}})\big)}{\sum_{\hat{t}=1}^{t-1} \alpha^{t-1-\hat{t}}}. \tag{1}$$

Because of the $\alpha$ parameter in the definition of $TDR$, for us, a history is a *timed-stamped* collection of past joint actions. This is very different from [27] which only uses the history to extract the distribution of joint actions and considers it to be a *collection (without timestamps)* of past joint actions. When $\alpha = 1$, the definitions of regret and time-discounted regret coincide.

Suppose $\Phi_c$ is the set of all functions from $A$ to $A$ that are *constant* functions, i.e., if $f$ is in $\Phi_c$, there must exist an action $a' \in A$ such that for all $a \in A$, $f(a) = a'$. The *time-discounted external regret* with respect to $\Phi_c$ is then simply given by

$$TDER_{i,\Phi_c}(t) = \max_{\hat{a} \in A} \frac{\sum_{\hat{t}=1}^{t-1} \alpha^{t-1-\hat{t}}\big(u_i\big(\hat{a}, a_{-i}^{\hat{t}}\big) - u_i(a^{\hat{t}})\big)}{\sum_{\hat{t}=1}^{t-1} \alpha^{t-1-\hat{t}}}.$$

In other words, $TDER_{i,\Phi_c}$ only considers constant functions when computing time-discounted regret. We define the time-discounted external regret with respect to action $\hat{a}$ as

$$TDER_i(\hat{a}, t) = \frac{\sum_{\hat{t}=1}^{t-1} \alpha^{t-1-\hat{t}}\big(u_i\big(\hat{a}, a_{-i}^{\hat{t}}\big) - u_i(a^{\hat{t}})\big)}{\sum_{\hat{t}=1}^{t-1} \alpha^{t-1-\hat{t}}}. \tag{2}$$

Intuitively, $TDER_i(\hat{a}, t)$ is the regret for player $i$ due to the fact that she/he did not use the strategy to always play the action $\hat{a}$ in the past. We assume that for a rational player, the greater the regret with respect to an action $\hat{a}$ in the past, the more likely it is that the player will play the action $\hat{a}$ in the future.[2]

*Example 3:* Let us reconsider example 2 with $\alpha = 0.9$. The time-discounted external regret for player 1 with respect to action $h$ in the year 2013 ($t = 4$) is $TDER_1(h, 4)$, shown at the bottom of this page.

Observe that the weights 0.81, 0.9, and 1.0 are the weights for years 2010, 2011, and 2012, respectively.

A player is *rational* if, for each time $t$, the player chooses the action that caused the maximum time-discounted external regret in the past. Thus, our rationality constraints require that $\forall t \in [T]\backslash\{1\}, \forall i \in [N], \forall \hat{a} \in A\backslash\{a_i^t\}$, the following condition holds[3]:

$$TDER_i(\hat{a}, t) \leq TDER_i\big(a_i^t, t\big) \tag{3}$$

or, equivalently

$$\sum_{\hat{t}=1}^{t-1} \alpha^{t-1-\hat{t}}\big(u_i\big(\hat{a}, a_{-i}^{\hat{t}}\big) - u_i\big(a_i^t, a_{-i}^{\hat{t}}\big)\big) \leq 0. \tag{4}$$

Since our goal is to infer the payoff function of players (based on regret maximization) instead of computing an equilibrium, we do not have any constraint at $t = 1$, and the action at $t = 1$ is from the ground truth of the data set. *Bounded Rationality:* As players in the real world are rarely 100% rational, we introduce a parameter $\epsilon \in [0, 1]$ that captures the degree of rationality. The closer $\epsilon$ is to 1, the more rational the player is, while the closer $\epsilon$ is to 0, the more irrational the player is. We replace (3) (which assumes complete rationality)

[2] In simple terms, if the player had great regret about not doing something in the past, especially the recent past, then he is more likely to do it in the future, especially in the near future.

[3] Note that since we set a constraint for each player, this means that each player is maximizing its regret knowing that the other players are simultaneously maximizing their regret.

$$TDER_1(h, 4) = \frac{0.81(u_1(h, m) - u_1(l, m)) + 0.9(u_1(h, l) - u_1(m, l)) + (u_1(h, m) - u_1(h, m))}{0.81 + 0.9 + 1.0}.$$

with the equation in the following, which allows weaker notions of rationality:

$$\epsilon \cdot TDER_i(\hat{a}, t) \leq TDER_i\left(a_i^t, t\right)$$

or equivalently

$$\sum_{\hat{t}=1}^{t-1} \alpha^{t-1-\hat{t}}\left(\epsilon \cdot u_i\left(\hat{a}, a_{-i}^{\hat{t}}\right) - u_i\left(a_i^t, a_{-i}^{\hat{t}}\right)\right) \leq 0. \tag{5}$$

As $\epsilon$ and $\alpha$ are constants, this equation is linear. Each $u_i(-)$ term is a variable in this constraint. Let $LC$ be the set of all linear constraints generated by (5). We demonstrate them in the next example.

*Example 4:* By considering only the last two years of the history in Example 2, we observe that $H^T$ is

| $[T]$ | year | player 1 | player 2 |
|---|---|---|---|
| 1 | 2011 | m | l |
| 2 | 2012 | h | m |

If $\alpha = 0.9$ and $\epsilon = 0.8$, the rationality constraints are

$$(0.8\, u_1(l,l) - u_1(h,l)) \leq 0(0.8\, u_1(m,l) - u_1(h,l)) \leq 0$$
$$(0.8\, u_2(m,l) - u_2(m,m)) \leq 0(0.8\, u_2(m,h) - u_2(m,m)) \leq 0.$$

The result below states that $LC$ is polynomial in size.

*Proposition 1:* The number of variables occurring in $LC$ is polynomial in the number of players $N$, the number of actions $M$, and in the size of the history $T$.

One problem with $LC$ is that it may have multiple solutions, some of which may be trivial. An example of a trivial solution is when the utility function returns the same value for each joint action for each player. For instance, the maximal entropy solution of $LC$ (the entropy function is applied to all variables of $LC$) assigns the same utility to all combinations of players and joint actions.

*Proposition 2:* Suppose $\mathcal{U} = \{u_1, \ldots, u_N\}$ is a maximal entropy solution for $LC$. Then, for all joint actions $a, a'$ and all players $i$, $j$, $u_i(a) = u_j(a')$.

The proofs of the above two propositions are in The Appendix. Hence, given the assumptions in this article, maximal entropy is not a very effective way of choosing a solution.

### C. Function Form Payoff Matrix With Time-Weighted History

PIE was motivated by our ongoing counter-terrorism research. We have applied PIE to a multiplayer game involving the terrorist group LeT (responsible for the 2008 Mumbai attacks) and the governments of Pakistan and India as players. In such scenarios, a time-weighted history of players' actions is a representation of state of the world (which is a history of players' actions) at the time the player decides to take a new action. Time-weighted history captures the idea that recent actions may be more relevant than older ones.

The payoffs in (5) are in tabular forms. Suppose we assume that a payoff function for player $i$ at time $t$ is *any* (linear or nonlinear) function $\pi_i : \mathbb{R}^{m+1} \mapsto \mathbb{R}$ of the time-weighted history at time $t$. A *time-weighed history*, $w_t$, at time $t$ is an $m$-dimensional vector defined as follows:

$$w_t = \frac{\sum_{i \in \{1\ldots t\}} \alpha^{t-i} H_i}{\sum_{i \in \{1\ldots t\}} \alpha^{t-i}}.$$

That is, $\pi_i(a, w_t)$ takes an action $a$ and a time weighted history $w_t$ as input and outputs a payoff value, specifying the payoff to player $i$ of playing $a$ at time $t$ with respect to $w_t$.

We assume that a player chooses an action at time $t$ that has highest payoff with respect to the state of the world at time $t-1$. Let $a$ be the action of player $i$ at time $t$, and thus

$$\pi_i(a, w_t) \geq \pi_i(a', w_t) \quad \forall a' \in A_i,\ i \in [N],\ t \in [T]. \tag{6}$$

One such constraint needs to be written for each player $i$ and each time $t$. *Note that these constraints may be nonlinear.* It is easy to see that (6) generalizes (5). Note that this system of inequalities is feasible as assigning identical payoffs for all actions always satisfies (6). However, this solution is trivial. Hence, it is important to choose a "robust" solution to the above system. For real-world problems, we require that the selected solution satisfies the following properties.

1) The family of functions to which our payoff functions $\pi_i$ belong should not have arbitrary complexity, i.e., our hypothesis space should not allow arbitrary payoff functions to avoid overfitting. On the other hand, we should allow somewhat complicated nonlinear payoff functions to avoid oversimplification.
2) While it is reasonable to assume that players react to game history and use actions which would generate high payoffs, we cannot assume that each and every player always adheres to this heuristic at all times. Therefore, our algorithm must admit the possibility that some points in the history may violate (6). However, (6) should hold for most of the game history.
3) Last but not least, there must be a tractable algorithm to select a solution of these constraints so that PIE can apply to real-world strategic games involving many players and dozens of actions. While the current best approach [27] in the literature has been applied to games of up to six players with three actions for each player, they do not discuss the runtime of their approach. Our experiments will show that our best approach is one to two orders of magnitude faster.

## IV. Solutions

In this section, we present three approaches to select a solution of the system of constraints defined in Section III. Note that the CBS and SCA approaches are designed for the $LC$ constraints in (5), while the SVMM is devised for the constraints in (6).

### A. CBS

The CBS uses $LC$ [see (5)]. The classical way to choose one solution of $LC$ is to choose the maximum entropy solution. However, as proved earlier in Proposition 2, this is not useful as the maximum entropy solution assigns the same utility to all combinations of players and joint actions. In order to avoid this, we choose a centroid-based approach. The *centroid solution* of $LC$ is the mean position of all points satisfying $LC$. Unfortunately, computing the centroid of a convex region is computationally very complex, and even approximating it is $\#P$-hard [36]. We, therefore, approximate the centroid by using hit-and-run (HAR) sampling [37]. In HAR sampling, we start with a randomly selected solution of the constraints (point in the polytope). We then randomly identify a direction and distance and head in that direction for the selected distance from the last sampled point. If we are still within the polytope, this becomes our next sampled point. If the new point is outside the polytope, we regenerate a distance and direction until a valid point within the polytope is found. This process is iterated until the desired number of sample points is generated. HAR sampling allows us to sample points from a convex polytope uniformly at random in time polynomial in the number of dimensions (the number of variables of $LC$). We approximate the centroid by taking the componentwise mean of the sampled payoffs.

*Proposition 3:* The centroid approximation described earlier is a solution of $LC$.

This proposition follows as the centroid approximation is a convex combination of the solutions of $LC$.

### B. SCA

In the SCA, we again use only $LC$ [see (5)] and allow the rationality constraints to be violated by introducing a slack variable in each constraint in $LC$. These slack variables are denoted $s_{i,a,t}$ in the revised linear program (RLP) given in the following. We then find a solution of $RLC$ that minimizes the sum of the slack variables which, in a sense, minimizes the amount of violation of the constraint. The RLP is shown in the following:

$$\min_{s,u} \sum_{i,a,t} s_{i,a,t}$$
$$\text{s.t}: \sum_{\hat{t}=1}^{t-1} \alpha^{t-1-\hat{t}} \left(\epsilon u_i\left(a, a_{-i}^{\hat{t}}\right) - u_i\left(a_i^t, a_{-i}^{\hat{t}}\right) + s_{i,a,t}\right) \leq 0 \quad \forall i \in [N],\ a \in A,\ t \in [T]. \tag{7}$$

The slack variables in (7) are inside the parentheses to normalize for the history length and time decay of payoffs.

### C. Payoff Inference Using SVMM

In this section, we present a novel approach that uses SVMs to find a "good" candidate solution to the system of inequalities given in (6). Here, we use a set $CONS$ of constraints, which are generated by (6). Each constraint generated by (6) has a left-hand side and a right-hand side. We encode the left- and right-hand sides of the inequalities in (6) as points in a space. If a point encodes the right-hand side of an inequality, it is assigned a label 1; otherwise, it is assigned a label 0. We then run the classification algorithm. The decision function for the learned classifier is the desired payoff function. We now describe this method in more detail.

*Encoding Points in Game History:* Let the number of actions of player $i$ be $n_i = |A_i|$. Let $a$ be an action of player $i$, whose index is equal to $\sum_{j\in\{1\ldots i-1\}} |A_j| + k$. Consider an $m+1$ tuple $(a, h)$, where $a \in A_i$ and $h$ is an $m$-dimensional point. Let $V : \mathbb{R}^{m+1} \mapsto \mathbb{R}^{(m+1)n_i}$ be a map that takes this $m+1$'th tuple as input and outputs an $(m+1) * n_i$-dimensional vector. Map $V$ is defined as follows:

$$V(a,h)[(k-1)*(m+1)+1] = 1$$
$$V(a,h)[(k-1)*(m+1)+1+j] = h[j] \quad \forall j \in \{1\ldots m\}.$$

All other entries of $V(a,h)$ are 0. Suppose player $i$ plays action $a$ at time $t$. Then, $V(a, w_{t,i}, w_{t-1,-i})$ is labeled 1. For all actions $a' \neq a$, $V(a', w_{t,i}, w_{t-1,-i})$ are labeled 0. The following example shows how this works.

*Example 5:* Consider a three-player game with players 1–3 with two actions (namely, actions 1 and 2) for each player. A point in the history of the game is represented by a 6-D binary vector, e.g., the vector (1,0,0,1,1,0) represents the fact that players 1–3 played actions 1, 2, and 1, respectively. Let the current state of the world be given by the vector $w_t = (w_1, w_2, w_3, w_4, w_5, w_6)$. Assume that at any time, player 1 plays a myopic best response to this state of the world. For simplicity, let payoffs be a linear function of the state of the world. The payoff for playing action 1 by player 1 is $p_1 = a_1 + \sum_{i\in\{1..6\}} a_{1i} w_i$. Similarly for action 2, $p_2 = a_2 + \sum_{i\in\{1..6\}} a_{2i} w_i$. Thus, the payoff function can be represented as a 14-D vector $p = (p_1, p_2)$. Furthermore, assume that player 1 actually chooses action 1 as the best response, then

$$p_1 >= p_2. \tag{8}$$

We now encode the RHS as $V(1, w_t) = (1, w_1, w_2, w_3, w_4, w_5, w_6, 0, 0, 0, 0, 0, 0, 0)$ and the LHS as $V(2, w_t) = (0, 0, 0, 0, 0, 0, 0, 1, w_1, w_2, w_3, w_4, w_5, w_6)$. If we use SVM to learn a separating hyperplane $W$ for points $V(1, w_t)$ and $V(2, w_t)$ such that $V(1, w_t)$ is on the positive side and $V(2, w_t)$ is on the negative side, then we have

$$W^T V(1, w_t) > 0, \quad W^T V(2, w_t) < 0. \tag{9}$$

Thus, we have $W^T V(1, s) > W^T V(2, s)$ and $W$ is a 14-D vector representing a feasible solution of (9).

Going back to the general case, let $\mathcal{E}$ be the function that takes a given game history as input and outputs a payoff value, the labeling, and encoding of points as defined earlier. We now describe the relationship between the SVM classifier applied to points given by mapping $V$ and the system of inequalities given by (6) with the help of the following two propositions—proofs are in The Appendix.

*Proposition 4:* The system of inequalities given in (6) is feasible for a game history $H$, i.e., we can find payoff functions such that all the inequalities are satisfied if the SVM algorithm can find a separator for encoding $\mathcal{E}(H)$.

*Proposition 5:* If the SVM algorithm can find a separator that misclassifies $n_1$ points with label 1 and $n_0$ points with label 0, then we can find payoff functions such that at most $n_0 + n_1$ of the inequalities in (6) are violated.

## V. Implementation and Experiments

We implemented CBS, SCA, and SVMM, as well as the ICEL algorithm [27]. Section V uses synthetic data (with known payoff functions to evaluate these algorithms' accuracy). Section VI uses the real-world MAROB data about ten terrorist groups [29] with actions by both the terrorist groups and the government of the country involved. Section VI-B also uses a very fine-grained counter-terrorism data set with three actors: the terror group LeT [2] which carried out the Mumbai attacks and the governments of Pakistan and India. Section VII compares our best algorithm with the ICEL algorithm. We used the MAROB data to compare ICEL and our SVMM method. Because of Proposition 2, we could not apply ICEL to the synthetic data, and because the LeT data contained a host of environmental variables, we could not apply ICEL to that either. For the experiments on both the synthetic and real data sets, the discount factor $\alpha$ is set to 0.9.

### A. Generation of Synthetic Data

We wrote R code to generate random games with random linear payoff functions and a random state of the world at each time. A payoff function is represented as a vector of coefficients of a linear function. Each of the payoff functions and the state of the world at each time point is a uniformly randomly directed positive vector of norm 1. After generating the payoffs and the state of the world at different times, an action history for all players is generated assuming best response. We are not simulating a game. Instead, each time point is a "what if" scenario, where each player is presented with a state of the world and they choose the best response as per their payoff functions. The code to generate random games varies the following inputs.

| | |
|---|---|
| np | Number of players |
| na | Number of actions for each player |
| n | Length of history for each game |
| noise | Probability an action is chosen uniformly at random |
| seed | The seed for random number generation |

The state of the world is thus an $(na * np)$-dimensional vector. Payoff for each action is a linear function of the state of the world, and hence, it is represented as an $(na * np)$-dimensional vector of coefficients. Thus, each player has $na$ such vectors. [4] We introduce noise into our experiments by allowing each player, at each time step, to play either a random response independently at random with probability given by parameter "noise" or a best response to the current state of the world. To evaluate quality of payoffs learned, we learn the $np * na^2$ length vector of parameters of each player's payoff function. We measure the quality of our three payoff learning algorithms by comparing this vector with the actual payoff function vectors using Pearson correlation coefficients (PCCs).

[4]For most experiments on synthetic data, we have $na = np = 3$. Thus, each payoff function is a 9-D vector. As there are three payoff functions per player (one per action), we are trying to learn a total of nine vectors, each of which is 9-D.

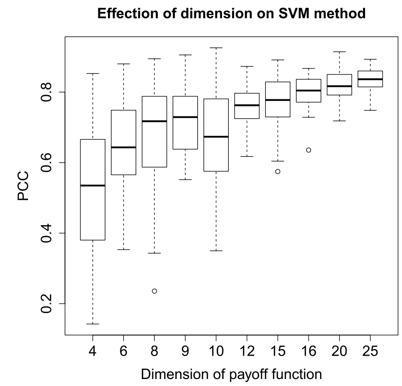


Fig. 1. Effect of payoff function dimension on the performance of the SVM method for synthetic data.

Sections V-B–V-F compare the three algorithms presented in this article in order to identify which one is best– both from an accuracy and from a run-time perspective.

### B. Performance of SVM-Based Method

We use a linear soft margin SVM classifier using the R interface to libsvm [38]. The hyperparameter for tuning this SVM is the cost of misclassification $C$. We tried the values of $C \in \{0.01, 0.1, 1, 10, 100\}$ and chose the best-forming SVM model. However, we also report the overall results (encompassing all five values of $C$). The choice of $C$ turns out to not be critical to the performance of our algorithms. SVMM performs very well with median PCC above 0.8 for games with five players and five actions for each player and median PCC between 0.6 and 0.8 for most of the smaller games. In addition, performance degrades slowly with noise. We now analyze SVMM's performance in more detail.

*1) Effect of Dimension of Payoff Function (SVMM):* Fig. 1 shows the effect of dimension of the payoff function on the performance of the SVMM.[5] Here, the number of sampled history points is 1000 and the noise parameter is set to 0. Surprisingly, the performance improves with dimension of the payoff function. This will be discussed in detail later in this section.

*2) Effect of Noise (SVMM):* Fig. 2 shows the effect of noise on SVMM's performance. The number of samples is 1000 and the dimension of the payoff function is 9. We note that performance degrades gracefully under noise. For zero noise, the median PCC value is 0.73, whereas even with noise as high as 0.3 (i.e., with probability 0.3, a player chooses to play a random action instead of the best response), we still get a median PCC of 0.66.

*3) Effect of Length of History (n) (SVMM):* Fig. 3 shows the effect of the length of the history on SVMM's performance. There is slight improvement in median PCC, as $n$ increases from 200 to 1000.

[5]The figure was plotted using the standard "boxplot" function in R (http://www.r-bloggers.com/boxplots-and-beyond-part-i/). The boxes denote the range of 25th and 75th percentiles. The line in the box is the median. The upper and lower lines outside the box are the "nominal" range of values and the circles are outliers.

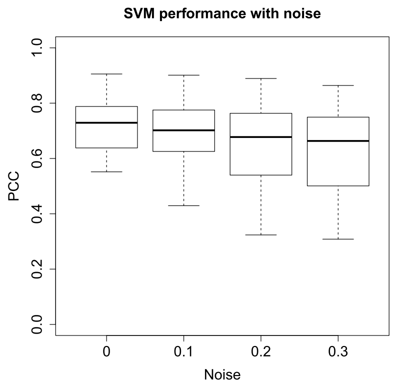


Fig. 2. Effect of noise on the performance of the SVM method for synthetic data.

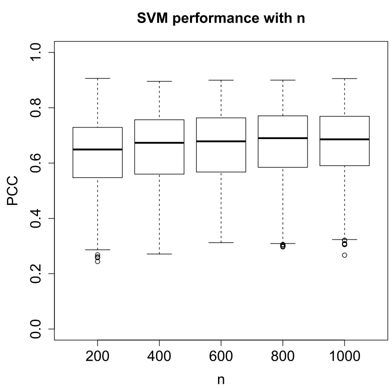


Fig. 3. Effect of history length on the performance of the SVM method for synthetic data.

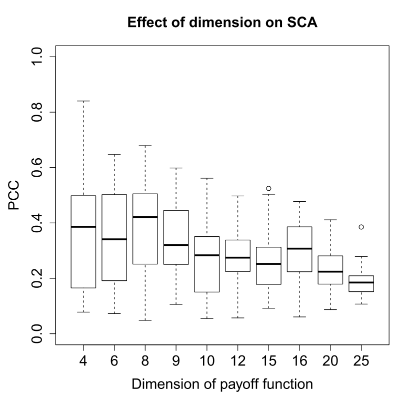


Fig. 4. Effect of dimension of payoff function on the performance of the SCA for synthetic data.

### C. Performance of SCA

In this section, we describe the performance of the SCA method and show that it is inferior to the SVMM.

*1) Effect of Dimension of Payoff Function (SCA):* Fig. 4 shows the effect of dimension of the payoff function on SCA's performance. The number of sample history points is 1000 and the noise parameter is set to 0. Performance degrades with dimension of the payoff function.

*2) Effect of Noise (SCA):* Fig. 5 shows the effect of noise on SCA's performance. The number of samples is 200 and

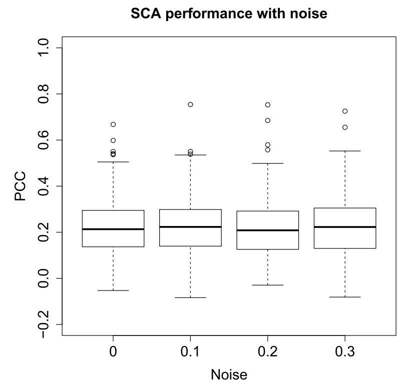


Fig. 5. Effect of noise on the performance of the SCA for synthetic data.

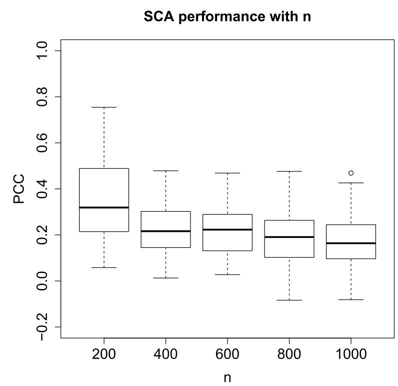


Fig. 6. Effect of history length on the performance of the SCA for synthetic data.

the dimension of the payoff function is 9. While performance does not degrade significantly with noise, overall performance is poor (overall PCC median of 0.22).

*3) Effect of History Length (SCA):* Fig. 6 shows the effect of history length on SCA's performance. Here, noise is 0 and dimension of the payoff function is 9. Somewhat counter intuitively, performance degrades with history length. This could be because the number of slack variables increases linearly with the length of history. Thus, the degrees of freedom of the model are potentially higher with a longer history, and thus, a longer history can lead to overfitting.

### D. Performance of Centroid-Based Method

In this section, we study CBS's performance and show that it is far inferior to the SVMM.

*1) Effect of $\epsilon$ (CBS):* Fig. 7 shows the effect of $\epsilon$ on CBS's performance. Here, the length of history is 200 and the dimension of the payoff function is 9. The noise is 0. The overall performance of CBS is better than that of SCA but much worse than the SVM-based method (PCC median of 0.45 for $\epsilon = 0.9$).

*2) Effect of Dimension of Payoff Function (CBS):* Fig. 8 shows the effect of dimension of the payoff function on the performance of CBS. Here, the number of sample history points is 1000 and the noise parameter is set to 0. Performance shows no discernible trend.

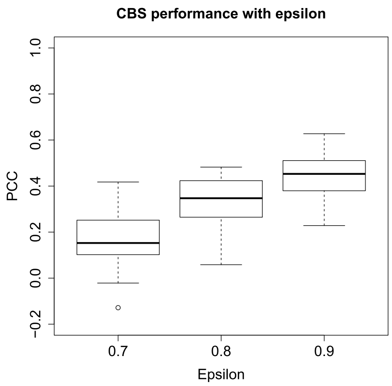


Fig. 7. Effect of epsilon on the performance of CBS for synthetic data.

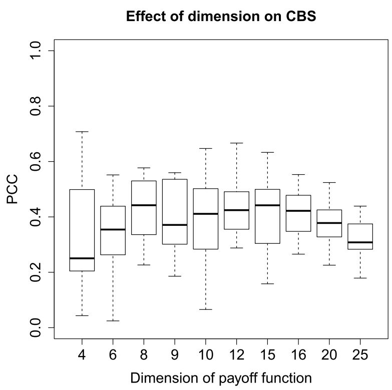


Fig. 8. Effect of dimension of payoff function on the performance of CBS for synthetic data.

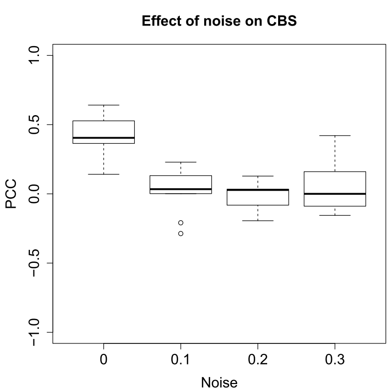


Fig. 9. Effect of noise on the performance of CBS for synthetic data.

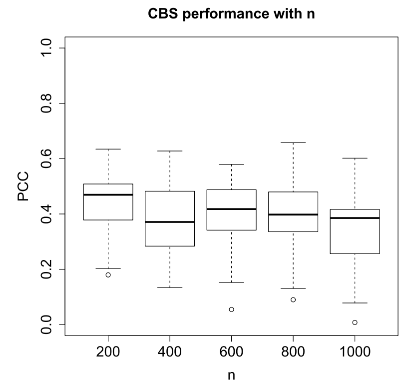


Fig. 10. Effect of history length on the performance of CBS for synthetic data.

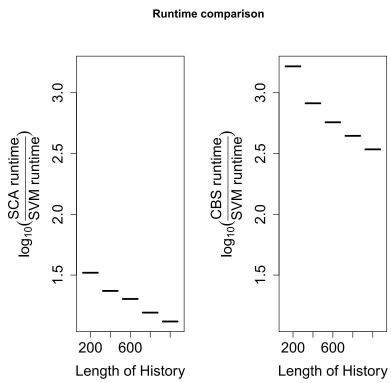


Fig. 11. Comparison of runtime of various methods for synthetic data.

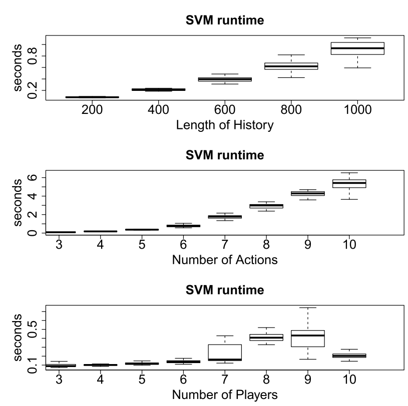


Fig. 12. Runtime of the SVM method.

*3) Effect of Noise:* Fig. 9 shows the effect of noise on the performance of CBS. Here, the number of samples is 1000 and the dimension of the payoff function is 9. Performance degrades sharply with noise and, even with 10% noise, is close to random.

*4) Effect of Length of History (CBS):* Fig. 10 shows the effect of history length on the performance of CBS. Here, noise is 0 and the dimension of the payoff function is 9. Performance does not improve with $n$ and shows no discernible trend.

## *E. Runtime Comparison of CBS, SCA, and SVMM*

Fig. 11 shows the relative runtime performance of the three methods for varying lengths of histories. The SVMM is faster than SCA by an order of magnitude and faster than CBS by two to three orders of magnitude. For this comparison, $na$ and $np$ are 3. The actual performance time of the SVMM for varying values of $n$, $na$, and $np$ is given in Fig. 12.

## *F. Discussion of the Results*

*1) Effect of Dimension of Payoff Function:* The effect of the dimension of the payoff functions on the performance of the SVMM, SCA, and CBS is depicted in Figs. 1, 4, and 8, respectively. We see that SVMM's performance improves with dimensionality. This is counter-intuitive as the performance of most classifiers degrades with dimension. However, for the payoff inference problem, the number of constraints and, hence, the number of points increase with the payoff function's dimension (for a fixed length of history). Thus, while the complexity of the classifier increases with data dimensionality, we have more data to learn from, and hence, performance improves with dimensionality. In the case of SCA, we see that performance degrades with dimensionality. For SCA, the number of slack variables is the product of history length and dimension of the payoff function. Thus, the increase in dimension leads to an increase in the number of the slack variables. We hypothesize that in SCA, this increase in the number of slack variables with increase in dimensionality leads to performance degradation. CBS shows no discernible trend with increasing dimensionality. First, we are only approximating the centroid. Second, the centroid is very sensitive to individual constraints. Thus, CBS chooses an approximation to a feasible representative solution that is very sensitive to individual constraints. Therefore, it is not surprising that its performance is erratic.

*2) Effect of Noise:* The effect of noise on the performance of the SVMM, SCA, and CBS algorithms is shown in Figs. 2, 5, and 9, respectively. SVMM's performance degrades gracefully with noise. As soft margin SVMs evolved from hard margin SVMs to handle misclassification, this graceful degradation is expected. SCA's performance remains more or less constant and very poor with and without noise. SCA accommodates noisy points by allowing constraints to be violated by allowing negative slack variables, and hence, some robustness to noise is expected. However, a total lack of trend is a bit surprising. As noted earlier, centroid is very sensitive to individual constraints, and hence, extreme sensitivity to noise, as shown in Fig. 9, is expected.

*3) Effect of History Length:* The effect of history length on the performance of SVMM, SCA, and CBS is shown in Figs. 3, 6, and 10, respectively. SVMM's performance improves slightly when $n$ increases. Thus, the SVMM learns a better classifier with more data. Surprisingly, for synthetic data, it seems that the SVMM is able to learn a very good classifier even with $n = 200$. SCA again shows the trend of degrading performance with the increasing number of constraints. Performance of CBS is again erratic.

*4) Runtime:* The runtimes of the SVMM, SCA, and CBS are compared in Fig. 11. CBS is easily the worst. SVMM's runtime increases with length of the history (see Fig. 12) because the problem size varies linearly with the length of the history. The corresponding increase with the number of actions is faster, as the problem size varies quadratically with the number of actions. SVMM's runtime increases with the number of players in general; however, when the number of players is ten, it suddenly drops. We do not have a good explanation for this behavior, hope to find one in the future work.

*Bottom Line:* We conclude by stating that of the three algorithms presented in this article, SVMM achieves significantly higher accuracy than SCA and CBS, is more robust to noise, and performs much better and faster than the other two methods. It gives excellent performance on relatively large games (median PCC is above 0.8 for games with five players, five actions per player).

## VI. Experiments on Real-World Data Sets

### *A. MAROB Experiments*

We ran tests on all ten terror groups in the MAROB [29] data set for which at least 20 rows of data are available. We aggregated low-level MAROB actions (both by the group and the government of the nation where the group is based) into high-level actions. The high-level actions involved two actions each for the group (political engagement with the government, militant activities) and for the government (political engagement with the group and suppression of the group). Each of these actions can be carried out at mild, medium, and intense extents. Hence, each player can take one of nine actions, leading to 81 total joint actions. We ran experiments with data about ten group/nation pairs.[6] As the ground truth about payoffs for the terrorist groups cannot be tested directly, in order to test the validity of the payoffs learned, we made predictions about actions carried out by terror groups and governments and checked the accuracy of these predictions. The mean number of actions for the government player denoted by $G$ is 4.1 and the mean for the terror organization denoted by $TO$ is 5.6. The mean number of joint actions (product of actions of $G$ and $TO$) is 25.30. The history length of each game is between 20 and 25.

There are two issues related to these data that must be mentioned.

1) We only have 20 data points for each group/nation pair.
2) We do not know the ground truth.

Therefore, we evaluate the quality of the learned payoffs as follows. We compute Spearman's rank correlation coefficient (SCC) between predicted payoffs and the binary vector representing actual actions performed during the time period. While the payoffs are reals in $[0, 1]$, we are correlating them with binary variables in $\{0, 1\}$, and thus even in the best case, we cannot expect the correlation to be 1. For example, for five actions, a point in history may be $(1, 0, 0, 0, 0)$. It will be correctly predicted by a payoff vector, such as $p = (1, p_2, p_3, p_4, p_5)$, $p_i < 1\ \forall i \in \{2 \ldots 5\}$. Assuming that $p_i \neq p_j\ \forall i, j \in \{1 \ldots 5\}$, the expected SCC for these data (assuming $p_i$ is uniformly distributed over $[0, 1)$) is 0.71 which is extremely good, considering the paucity of data. Fig. 13 gives the best expected SCC as a function of the number of actions of a player. We normalized the SCC between predicted payoffs and the actual binary action vector to arrive at normalized SCC (NSCC) based on the number of actions of a player.

*Comparison of Performance of the Three Methods (Marob):* Fig. 14 compares CBS, SCA, and SVMM. On this data set, SVMM uses the radial basis function as the kernel and the model is selected based on leave-one-out cross validation. Here again, SVMM performs well (median NSCC = 0.7) and comfortably outscores SCA (median NSCC = 0.6) and CBS (median NSCC = 0). While for some part of the data set CBS does well with 75th percentile NSCC value close to 0.4, the average CBS performs very poorly with a median NSCC of

[6] We pair group and nations based on the simplified assumption that groups have very limited interactions. While there might be a few exceptions, this is true in most real-world cases.

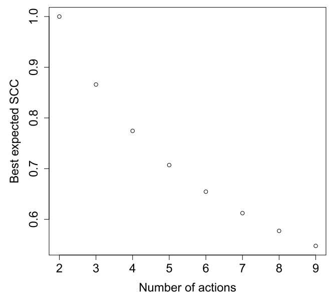


Fig. 13. Best expected SCC.

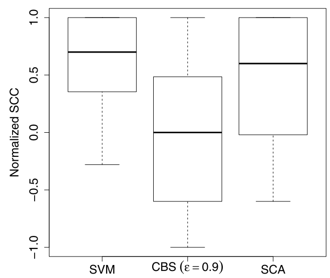


Fig. 14. Comparative performance of three methods on the MAROB data set.

close to 0, indicating near random performance. SCA performs well but not as well as the SVMM.

As in the case of synthetic data (see Fig. 9), CBS is very sensitive to noise. SVMM and SCA perform well because they both allow some constraints to be violated. As shown in Figs. 4 and 6, SCA's performance degrades with the number of constraints. However, since all the histories in the MAROB data set are short (∼20), SCA performs quite well. SVMM performs even better than SCA. We hypothesize that this is because SVMM accommodates nonlinear payoff functions using kernel methods and SCA allows only linear payoffs. This may be because real players' payoffs functions are not necessarily linear.

## B. LeT Experiments

We conducted extensive tests on the LeT data set [2], which contains 252 rows (months) of data about 700+ variables.[7] We choose 24 variables, which we considered relevant to our problem. These variables include six variables for various types of attacks carried out by LeT and eight for actions taken by the Pakistani government and military. The other ten variables, such as the existence of an international ban, the existence of conflict within LeT, split in LeT, and so on, are treated as *environmental variables*.[8]

The LeT data set models a relatively big game. Players can take many actions simultaneously. If we encode each combination of actions as a separate action, we will end up with 64 actions for LeT and 256 actions for Pakistani government.[9] This leads to a very high-dimensional encoding for this data set. As an illustration, for a given history, $H$ of this game, $\mathcal{E}(H)$ for Pakistani government would be $256*(1+336) = 86\,272$ dimensions (ignoring the environment variables). We now describe how we tackle the dimensionality problem.

*1) Independent Payoffs for Simultaneous Actions:* One natural way to reduce dimensionality is to assume that payoffs for simultaneous actions are independent and additive. However, different actions may require different levels of effort for the player and it is reasonable to assume that payoffs are proportional to effort. For example, if attacking a security installation requires double the effort of attacking civilian transport, then the payoffs for the two actions are comparable only if the payoff from attacking the security installation is double the payoff from attacking civilian transport. This is because the capability and resources of an organization are limited, and thus, to maximize the payoff, effort should be spent on actions that give maximum payoff for each unit of effort. Therefore, for this approach, we need to assign effort-based weights to players' actions. However, there is no reliable way of knowing how much effort was needed for each action of the player. Therefore, we rejected this approach.

Instead, we relax the constraints in (6) by assuming that regret for each action actually played at time $t$ is greater than or equal to regret for actions not played at time $t$. For example, assume that at time $t$, a player played actions $(0, 1, 0, 1, 0, 0)$ indicating that they took actions 2 and 4 out of possible actions in $\{1, \ldots, 6\}$. We then assume that regrets for actions 2 and 4 were higher than regrets for other actions at time $t$. The other alternative could have been encoding each of the possible combinations of actions as a separate action, which leads to 64 possible distinct actions at each time step.

*2) Evaluation of Quality of Learned Payoffs:* We evaluate the quality of prediction in two ways. First, we compute the SCC of events and payoffs for each time period from the test data. We compute SCC between predicted payoffs for actions and the binary vector representing actual occurrence of the events during the time period. We use this method to compare the performance of SVMM, CBS, and SCA. Second, for SVMM, we compute the quality of predictions using area under the RoC curve as our metric. We do not use predictions to evaluate CBS and SCA methods because these methods do not extend naturally to prediction, and prediction is not our primary objective.

*3) Comparison of the Methods:* Fig. 15 presents comparative performance of the three methods on the LeT data set. Again, SVMM performs well and clearly outperforms SCA and CBS. The NSCC for SVMM when the player considered is LeT in the data set is 0.59. The NSCC for SVMM when the player considered is Pakistan is 0.80. The predictive performance of SVMM is good with area under single point RoC curve of 0.74 and 0.85 for LeT and Pakistan, respectively (see Fig. 16).

[7]It includes details about attacks carried out by LeT, communications campaigns, and rallies organized by LeT. It also includes actions by the state (Pakistan) and international actors (U.S., India, EU, and so on), such as arrests, tribunals, and killings related to members of LeT.

[8]Environmental variables can be seen to be actions of another player (similar to the *chance player* in classical game theory), whose actions we cannot predict (or are not interested in predicting). Nevertheless, these actions do have an effect on payoffs of other players.

[9]The data set does not attribute other actions to specific third part players, such as the U.S. or India.

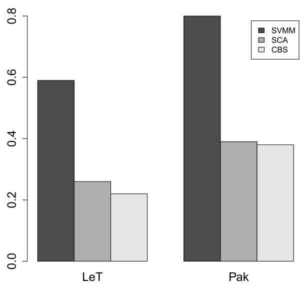


Fig. 15. Performance of SVMM, SCA, and CBS on LeT data set.

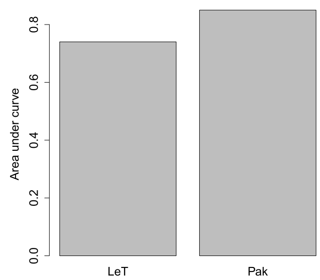


Fig. 16. Predictive performance of SVMM on LeT data set.

The performance of the SVMM is much better than the other two methods. We think that the reasons are threefold. First, the SVMM is robust to noise. Second, the SVMM allows for nonlinear payoff functions. Third, SVMM performs better with more constraints and higher dimensional payoff functions.

*4) Policy Options Based on Payoffs Learned About LeT:* We used the payoffs learned for different actions to qualitatively estimate the values of the different actions that the players in the LeT case study can perform. In particular, the payoffs suggest the value or lack of value of certain actions by the two players (LeT and the Government of Pakistan), as well as actors who can reshape the environment surrounding LeT, such as India or the U.S.

*a) Actions by Pakistan:* The payoffs we inferred using PIE have made some concrete discoveries. First, they suggest that arrests of LeT personnel by Pakistan are ineffective—a finding consistent with [2] where the authors found that arrests of LeT personnel can actually be followed by attacks on soft targets. Second, they suggest that release of operatives from prison is also not effective in curbing attacks by LeT. This is consistent with Rule (PST-4) in [2], which says that LeT attacks symbolic sites three months after months when the Pakistani government releases 0–9 LeT prisoners and LeT has locations across the border in India. It is also consistent with Rule (AA-4) in [2], which says that LeT attacks carriy out attempted attacks one month after 0–2 LeT personnels are released.

PIE has additionally shown that a ban on LeT by the Pakistani government can have an effect on curbing attacks on transportation sites, public sites, and tourist sites. This is a new result, consistent with results in [2], which shows that having a ban is linked to LeT backed attacks on Hindus and not having a ban is linked to attacks on professional security forces, security installations, and armed clashes.

Another interesting new finding by PIE is that a freeze of LeT assets by Pakistan are linked to a much higher payoff for targeting professional security forces, security installations, and Hindus.

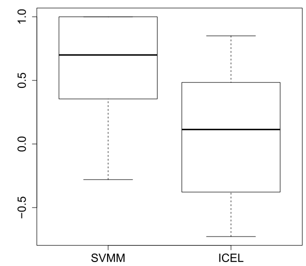


Fig. 17. Comparative performance of ICEL and SVMM.

Finally, PIE's payoffs confirm conventional wisdom that military support by Pakistan leads to higher payoffs for targeting security forces, installations, and public structures.

*b) Actions by international actors (primarily India and the U.S.):* We know from [2] that arrests of LeT personnel are linked to a reduction in attacks on hard targets (such as security installations) but are linked to an increase in attacks on softer targets—PIE's inferred payoffs confirm this by showing that such arrests lead to increased attacks on civilians on the basis of ethnicity (Hindus) and tourist sites.

Likewise, the imposition of international bans can lead to more attacks on civilian targets and fewer attacks on security installations, suggesting that external pressure on LeT causes them to move from attacking hard targets to softer ones.

A new result derived by PIE is that asset freezes reduce the payoffs for LeT to carry out attacks on civilian targets.

## VII. Comparison With ICEL

We compare the performance of our best algorithm (SVMM) against the ICEL algorithm of [27]. ICEL assumes that players play a correlated equilibrium. The input to ICEL is the game history and corresponding outcomes, which depend on joint actions of the players. The output is a joint distribution over the player's actions. The learned distribution is a correlated equilibrium for the corresponding inferred payoffs.

We evaluated the performance of ICEL against SVMM using the NSCC metric on the MAROB data set. We could not evaluate ICEL on the LeT data set because it has environmental variables in addition to player actions and ICEL is not applicable to games with environmental variables. If the payoffs learned by ICEL correspond to a correlated equilibrium actually played by the players, we can expect good rank correlation between chosen actions and payoffs. However, as can be seen from Fig. 17, this is not the case. The median NSCC is 0.1140 over all games (which suggests that ICEL is only marginally better than random noise). NSCC is above 0.5 in only 3 out of 20 instances (ten games, two players per game). In comparison, SVMM exhibits far superior performance with median NSCC of 0.7.

We believe that the poor performance of ICEL stems from two factors. First, we have no knowledge of the outcomes, only of game history. We empirically observed that in the absence of any information about the outcomes, the ICEL convex

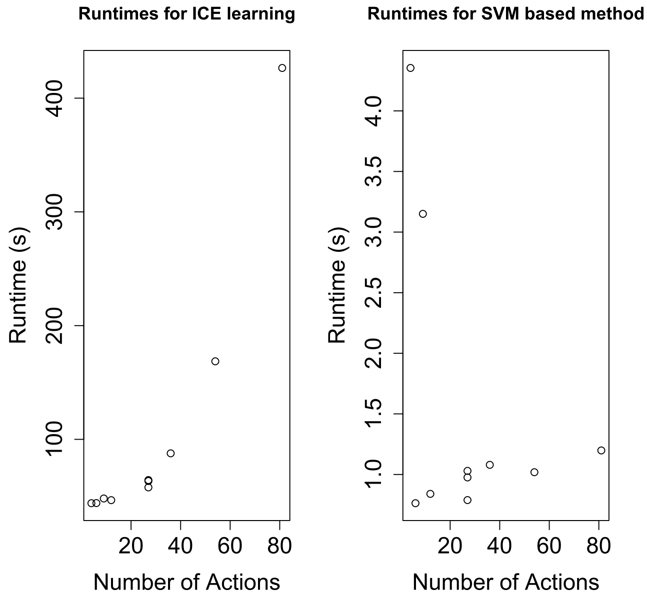


Fig. 18. Comparative runtime for ICEL and SVMM.

program simply converges to the distribution of actions actually played by the players (mean Kullback–Leibler divergence between actual and learned distributions is 0.043, max 0.088). Thus, the method corresponds to predicting that whatever happened in past will happen in the future with no notion of recency and dynamics. Second, in our real-world data, the players may not play a correlated equilibrium and our proposed dynamics, which are based on regret minimization and recency, may be a closer approximation of reality.

The runtime comparison is shown in Fig. 18. Here again, SVMM is one to two orders of magnitude faster than ICEL. However, this is a bit of apples and oranges comparison as the SVMM code is in R with a libsvm backend and the ICEL code uses the Python code provided by authors publicly at their website. We note that ability to use highly optimized, stable, and mature libraries provided by machine learning community for classification tasks is one of the advantages of our approach over other extant approaches.

## VIII. Conclusion

In this article, we have developed, for the first time, a method to infer payoffs for real-world games, under much more reasonable assumptions than past work. Specifically, unlike most past work, PIE is applicable to multiplayer games, allows players to not be fully rational, does not assume a coordination mechanism, does not assume a symmetric game, and is scalable, while other works are lacking in at least one of these aspects. Moreover, we apply our theory to real-world strategic games with two real-world counter-terrorism data sets based on two widely influential previous studies [2], [39]. Such individuals seek understanding—and our goal in learning these payoffs were to facilitate explaining the payoffs to such senior decision makers—rather than prediction. Toward this end, we propose three heuristics that may be used to learn payoffs of players in multiplayer real-world games, including one that builds upon SVMs—a tested technique in data mining that has never been used before for learning payoffs. Though the goal of this article is not prediction, we test our methods in three ways. We use a synthetic data set and a well-known terrorism data set [29] to see how well we can predict known payoff functions (synthetic data) and actions (MAROB data). Even though we have small amounts of data in both cases, they are bigger than those in the previous studies, and our best algorithm (SVMM) achieves good correlations. Our third test looks at ten years of data about the terror group LeT (responsible for the 2008 Mumbai attacks). We show that SVMM is both faster and much more accurate than ICEL [27], one of the best prior algorithms in the literature. Much work remains to be done. Even though PIE is more scalable than past work, there is still a long way to go. Moreover, explaining learned payoff functions to real-world decision makers also has many challenging aspects that deserve much more future study.

TABLE I
Summary of Notations

| Symbol | Brief synopsis |
|---|---|
| $[N]$ | The set of players |
| $\mathcal{A}$ | The set of all possible joint actions |
| $A_i$ | The set of actions of Player $i$ |
| $u_i$ | The payoff matrix for Player $i$ |
| $\mathcal{U}$ | The set of payoff matrices for the game |
| $[T]$ | The set of past time points |
| $a$ | The vector representing joint actions of all players |
| $(k, a_{-i})$ | The vector $a$ with action of player $i$ modified to action $k$ |
| $a^t$ | The vector of joint actions of all players at time $t$ |
| $H^t$ | The complete history sequence $< a^1, \ldots, a^t >$ till time $t$ |
| $w_t$ | The weighted time history till time $t$ |
| $\alpha$ | Time discount factor for reward |
| $\Phi$ | Family of modification functions for computing regret |
| $TDR_{i,\Phi}(t)$ | Time discounted regret for Player $i$ at time $t$ |
| $TDER_i(t)$ | Time discounted external regret for Player $i$ at time $t$ |
| $\pi_i(k, w_t)$ | Payoff function for Player $i$, (possibly non-linear). |

## Appendix

### A. Summary of Notations

Table I shows a summary of notations in this article.

### B. Additional Experimental Results

*1) Effect of Cost (SVM):* Fig. 19 shows the effect of the hyperparameter $C$ on the performance of the SVMM. We present this figure in order to show that for the synthetic data, the performance is quite stable and not very dependent on the choice of the hyperparameter $C$. Hence, tuning the SVM is straightforward.

*2) Actual Runtime for CBS and SCA Methods:* Figs. 20 and 21 show the actual runtime of CBS and SCA under various settings.

### C. Proof of Proposition 1

*Proof 1:* The number of constraints for each player at time $t \in [T]$ is $M - 1$. Each constraint has at most $T$ variables. Thus, the total number of variables that can occur in $LC$ is at most $(M - 1)NT$.

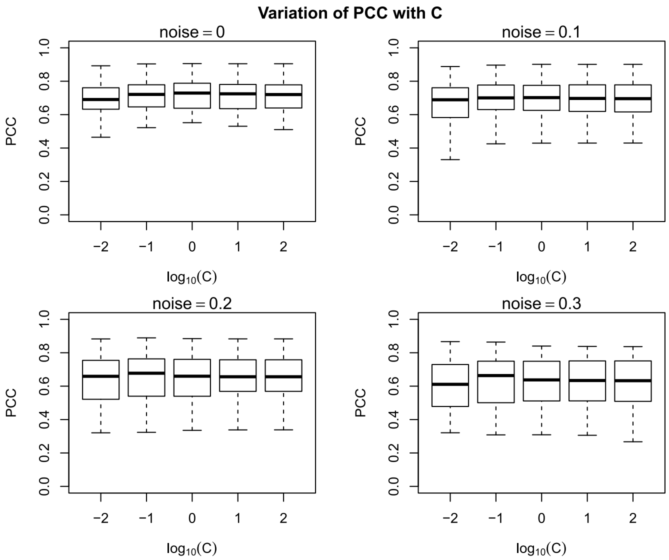


Fig. 19. Effect of cost on the performance of the SVM method for synthetic data.

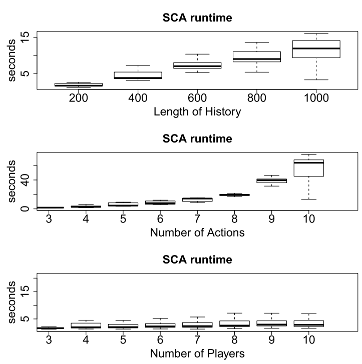


Fig. 20. Runtimes for SCA.

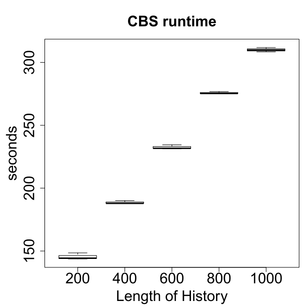


Fig. 21. Runtimes for CBS.

## D. Proof of Proposition 2

*Proof 2:* Let the entropy function be defined as follows:

$$-\sum_{i\in[N],a\in\mathcal{A}} u_i(a)\, ln(u_i(a)).$$

We obtain the maximum value of this function when $\forall i \in [N]$, $\forall a \in \mathcal{A}$ we can deduce that $u_i(a) = e^{-1}$. Since we know that when $\epsilon \in [0, 1]$

$$\sum_{\hat{t}=1}^{t-1} \alpha^{t-1-\hat{t}} \epsilon \cdot e^{-1} \le \sum_{\hat{t}=1}^{t-1} \alpha^{t-1-\hat{t}} e^{-1}$$

our rationality constraints in (5) are satisfied. For these reasons, it follows that the theorem holds. □

## E. Proof of Proposition 4

*Proof 3:* We note that for points on one side of decision surface, the value of decision surface is less than 0 and for the other side, it is greater than 0. Therefore, if points encoded for the RHS are on the positive side and LHS on the negative side, (6) is satisfied. Otherwise, we can flip sign of the decision function and achieve the same result.

## F. Proof of Proposition 5

*Proof 4:* Without loss of generality, we assume that points encoding the RHS of (6) are assigned positive labels and points encoding the LHS are assigned negative labels. A misclassified point LHS point can be assigned a decision value higher than the RHS point that can lead to at most one violated constraint. Similarly, a misclassified RHS point can lead to at most one violated constraint. Thus in all, we can have at most $n_0 + n_1$ violated constraints.

## REFERENCES


[1] T. C. Schelling, *The Strategy of Conflict*. London, U.K.: Oxford Univ. Press, 1960.

[2] V. Subrahmanian, A. Mannes, A. Sliva, J. Shakarian, and J. P. Dickerson, *Computational Analysis of Terrorist Groups: Lashkar-e-Taiba: Lashkar-e-Taiba*. Basel, Switzerland: Springer, 2012.

[3] J. P. Dickerson, A. Mannes, and V. Subrahmanian, "Dealing with lashkar-e-taiba: A multi-player game-theoretic perspective," in *Proc. Eur. Intell. Secur. Inform. Conf.*, Sep. 2011, pp. 354–359.

[4] D. P. Foster and R. V. Vohra, "Calibrated learning and correlated equilibrium," *Games Econ. Behav.*, vol. 21, no. 1, pp. 40–55, 1997.

[5] F. Xia, B. Jedari, L. T. Yang, J. Ma, and R. Huang, "A signaling game for uncertain data delivery in selfish mobile social networks," *IEEE Trans. Comput. Social Syst.*, vol. 3, no. 2, pp. 100–112, Jun. 2016.

[6] H. E. Lapan and T. Sandler, "Terrorism and signalling," *Eur. J. Political Econ.*, vol. 9, no. 3, pp. 383–397, 1993.

[7] W. Enders and T. Sandler, "Terrorism: Theory and applications," in *Handbook of Defense Economics*, vol. 1. Amsterdam, The Netherlands: Elsevier, 1995, pp. 213–249.

[8] M. Tambe, *Security Game Theory: Algorithms, Deployed Systerm, Lessons Learned*. Cambridge, U.K.: Cambridge Univ. Press, 2011.

[9] R. Lindelauf, P. Borm, and H. Hamers, "The influence of secrecy on the communication structure of covert networks," *Social Netw.*, vol. 31, no. 2, pp. 126–137, 2009.

[10] R. Lindelauf, H. Hamers, and B. Husslage, "Cooperative game theoretic centrality analysis of terrorist networks: The cases of Jemaah Islamiyah and al Qaeda," *Eur. J. Oper. Res.*, vol. 229, no. 1, pp. 230–238, 2013.

[11] T. P. Michalak, T. Rahwan, O. Skibski, and M. Wooldridge, "Defeating terrorist networks with game theory," *IEEE Intell. Syst.*, vol. 30, no. 1, pp. 53–61, Jan./Feb. 2015.

[12] Z. Wang, Y. Yin, and B. An, "Computing optimal monitoring strategy for detecting terrorist plots," in *Proc. 30th AAAI Conf. Artif. Intell. (AAAI)*, Feb. 2016, pp. 637–643.

[13] J. P. Dickerson, A. Sawant, M. T. Hajiaghayi, and V. Subrahmanian, "Preve: A policy recommendation engine based on vector equilibria applied to reducing let's attacks," in *Proc. IEEE/ACM Int. Conf. Adv. Social Netw. Anal. Mining*, 2013, pp. 1084–1091.

[14] A. Petrin, "Quantifying the benefits of new products: The case of the minivan," Nat. Bureau Econ. Res., Cambridge, MA, USA, Tech. Rep., 2001.

[15] A. Nevo, "Measuring market power in the ready-to-eat cereal industry," *Econometrica*, vol. 69, no. 2, pp. 307–342, 2001.

[16] J. Suzuki, "Land use regulation as a barrier to entry: Evidence from the texas lodging industry," *Int. Econ. Rev.*, vol. 54, no. 2, pp. 495–523, 2013.

[17] M. Conlin and V. Kadiyali, "Entry-deterring capacity in the texas lodging industry," *J. Econ. Manage. Strategy*, vol. 15, no. 1, pp. 167–185, 2006.

[18] A. Y. Ng and S. J. Russell, “Algorithms for inverse reinforcement learning,” in *Proc. ICML*, 2000, pp. 663–670.

[19] B. D. Ziebart, A. L. Maas, J. A. Bagnell, and A. K. Dey, “Maximum entropy inverse reinforcement learning,” in *Proc. AAAI*, 2008, pp. 1433–1438.

[20] G. Neu and C. Szepesvari, “Apprenticeship learning using inverse reinforcement learning and gradient methods,” Jun. 2012, *arXiv:1206.5264*. [Online]. Available: https://arxiv.org/abs/1206.5264

[21] S. Natarajan, G. Kunapuli, K. Judah, P. Tadepalli, K. Kersting, and J. Shavlik, “Multi-agent inverse reinforcement learning,” in *Proc. 9th Int. Conf. Mach. Learn. Appl. (ICMLA)*, Dec. 2010, pp. 395–400.

[22] D. Hadfield-Menell, S. J. Russell, P. Abbeel, and A. Dragan, “Cooperative inverse reinforcement learning,” in *Proc. Adv. Neural Inf. Process. Syst.*, 2016, pp. 3909–3917.

[23] T. S. Reddy, V. Gopikrishna, G. Zaruba, and M. Huber, “Inverse reinforcement learning for decentralized non-cooperative multiagent systems,” in *Proc. IEEE Int. Conf. Syst., Man, (SMC)*, Oct. 2012, pp. 1930–1935.

[24] B. Wiedenbeck and M. P. Wellman, “Learning payoffs in large symmetric games,” in *Proc. Int. Conf. Auto. Agents Multiagent Syst.*, May 2015, pp. 1881–1882.

[25] X. Wang and D. Klabjan, “Competitive multi-agent inverse reinforcement learning with sub-optimal demonstrations,” Jan. 2018, *arXiv:1801.02124*. [Online]. Available: https://arxiv.org/abs/1801.02124

[26] X. Lin, P. A. Beling, and R. Cogill, “Multiagent inverse reinforcement learning for two-person zero-sum games,” *IEEE Trans. Games*, vol. 10, no. 1, pp. 56–68, Mar. 2018.

[27] K. Waugh, D. Bagnell, and B. D. Ziebart, “Computational rationalization: The inverse equilibrium problem,” in *Proc. ICML*, 2011, pp. 1169–1176.

[28] C. Cortes and V. Vapnik, “Support-vector networks,” *Mach. Learn.*, vol. 20, no. 3, pp. 273–297, 1995.

[29] A. P. Asal, Victor, and J. Wilkenfeld. (Sep. 2008). *Minorities at risk organizational behavior data and codebook version 9/2008*. [Online]. Available: http://www.cidcm.umd.edu/mar/data.asp

[30] W. C. Stirling, “Conditional coordination games on cyclic social influence networks,” *IEEE Trans. Computat. Social Syst.*, vol. 6, no. 2, pp. 250–267, Apr. 2019.

[31] M. Soltanolkottabi, D. Ben-Arieh, and C.-H. Wu, “Modeling behavioral response to vaccination using public goods game,” *IEEE Trans. Computat. Social Syst.*, vol. 6, no. 2, pp. 268–276, Apr. 2019.

[32] I. Nevo and I. Erev, “On surprise, change, and the effect of recent outcomes,” *Frontiers Psychol.*, vol. 3, p. 24, Feb. 2012.

[33] C. G. Chorus, “A new model of random regret minimization,” in *Proc. EJTIR*, vol. 2, 2010, p. 10.

[34] I. Simonson, “The influence of anticipating regret and responsibility on purchase decisions,” *J. Consum. Res.*, vol. 19, pp. 105–118, Jun. 1992.

[35] M. Zeelenberg and R. Pieters, “A theory of regret regulation 1.0,” *J. Consum. Psychol.*, vol. 17, no. 1, pp. 3–18, 2007.

[36] L. A. Rademacher, “Approximating the centroid is hard,” in *Proc. SOCG*. ACM, 2007, pp. 302–305.

[37] D. Bertsimas and S. Vempala, “Solving convex programs by random walks,” *J. ACM*, vol. 51, no. 4, pp. 540–556, 2004.

[38] C.-C. Chang and C.-J. Lin, “Libsvm: A library for support vector machines,” *ACM Trans. Intell. Syst. Technol. (TIST)*, vol. 2, no. 3, p. 27, 2011.

[39] V. Subrahmanian, A. Mannes, A. Roul, and R. Raghavan, *Indian Mujahideen: Computational Analysis and Public Policy*. Springer, 2013.


**Haipeng Chen** received the B.S. degree from the University of Science and Technology of China, Hefei, China, in 2013, and the Ph.D. degree from Nanyang Technological University, Singapore, in 2018. He currently holds a postdoctoral position with the Department of Computer Science, Dartmouth College, Hanover, NH, USA. He has published various articles in top conferences and top journal in the field of artificial intelligence. His research lies in the general areas of artificial intelligence (AI), including machine learning, data mining, multiagent systems, and computational game theory and economics. Besides the theoretical foundations of AI, he is also interested at their intersections with various application domains, such as cybersecurity and intelligent transportation.

He was a member of the winning team for the Microsoft Malmo Collaborative AI Challenge and the runner-up team for the Innovation Demonstration Award of IJCAI’19. He has published various articles in top conferences, such as Proceedings of International Conference on Association for the Advancement of Artificial Intelligence (AAAI), Proceedings of International Joint Conference on Artificial Intelligence (IJCAI), Proceedings of ACM International Conference on Knowledge Discovery and Data Mining (KDD), and Proceedings of IEEE International Conference on Data Mining (ICDM), and top journals, such as the IEEE JOURNAL ON SELECTED AREAS IN COMMUNICATIONS, *ACM Transactions on Intelligent Systems and Technology*, and *Transportation Research Part C: Emerging Technologies*. He is a Program Committee Member for top conferences, such as AAAI, IJCAI, and Proceedings of International Conference on Autonomous Agents and Multiagent Systems (AAMAS), and a Journal Reviewer for top journals, such as the *European Journal of Operational Research*, the *IEEE Intelligent Systems* (IEEE-IS), the *Journal of Computer and System Sciences*, and *ACM Temperature-Controlled Logistics*.

**Mohammad T. Hajiaghayi** (F’20) received the Ph.D. degree from MIT, Cambridge, MA, USA, in 2005.

He is a Jack and Rita G. Minker Professor of computer science with the University of Maryland at College Park, College Park, MD, USA. He was a Postdoctoral Fellow at the School of Computer Science, Carnegie Mellon University, Pittsburgh, PA, USA, and in the Computer Science and Artificial Intelligence Laboratory, MIT. He has been a Distinguished Scientist in top labs in USA, such as AT&T Labs Research, Murray Hill, NJ, USA, Amazon Labs New York, NY, USA, Google Research, New York, Microsoft Research, New York, and IBM Research, Hawthorne, NY, USA. His main area of research is designing algorithmic frameworks. In particular, he designs efficient online algorithms, game theory algorithms, approximation algorithms, fixed-parameter algorithms, and streaming/parallel algorithms for big and complex networks and games, and often tests them in real-world and practical situations.

Dr. Hajiaghayi is a Guggenheim Fellow and an ACM Fellow.

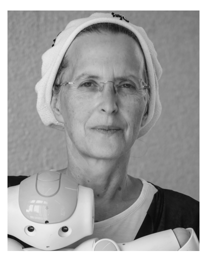

**Sarit Kraus** received the Ph.D. degree in computer science from the Hebrew University of Jerusalem, Jerusalem, Israel, in 1989.

She is currently a Professor of computer science with Bar-Ilan University, Ramat Gan, Israel. Her research is focused on intelligent agents and multi-agent systems (including people and robots). She has published over 400 articles in leading journals and major conferences, and has co-authored five books.

Dr. Kraus was an ACM, Proceedings of International Conference on Association for the Advancement of Artificial Intelligence (AAAI), and ECCAI Fellow. She is a member of the Board of Directors of the International Foundation for Multi-agent Systems (IFAAMAS). She has received many prestigious awards for her work. She was awarded the Proceedings of International Joint Conference on Artificial Intelligence (IJCAI) Computers and Thought Award, the ACM SIGART Agents Research Award, and the EMET Prize, and she was twice the winner of the IFAAMAS Influential Paper Award. She was a recipient of the advanced ERC grant. She was the IJCAI 2019 Program Chair.

**Anshul Sawant** studied computer science at IIT Delhi, New Delhi, India, and received the Ph.D. degree in computer science from the University of Maryland, College Park, MD, USA, in 2016, where he studied behavior of intelligent agents under the guidance of Prof. V. S. Subrahmanian and Prof. M. Hajiaghayi.

He is currently with Google, Mountain View, CA, USA, where he works on recommender systems for Google shopping.

**Edoardo Serra** received the Ph.D. degree in computer science engineering from the University of Calabria, Rende, Italy, in 2012.

He was a Visiting Researcher in the Computer Science Department, University of California at Los Angeles, Los Angeles, CA, USA, from October 2010 to July 2011. After the Ph.D. degree, he held a postdoctoral position at the University of Calabria in 2012, and he was a Research Associate at the University of Maryland, College Park, MD, USA, from 2013 to August 2015. Since August 2015, he has been an Assistant Professor with the Computer Science Department, Boise State University (BSU), Boise, ID, USA. His research interest is in the field of AI with applications in cybersecurity and national security.

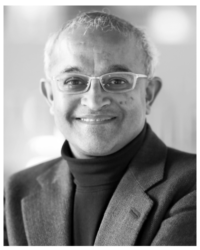

**V. S. Subrahmanian** is a Distinguished Professor in cybersecurity, technology, and society and the Director of the Institute for Security, Technology, and Society, Dartmouth College, Hanover, NH, USA. He previously served as a Professor of computer science at the University of Maryland, College Park, MD, USA, from 1989 to 2017, where he created and headed both the Laboratory for Computational Cultural Dynamics and the Center for Digital International Government. He also served over six years as the Director for the Institute for Advanced Computer Studies, University of Maryland. He is an expert on big data analytics, including methods to analyze text/geospatial/relational/social network data, learn behavioral models from the data, forecast actions, and influence behaviors with applications to cybersecurity and counterterrorism. He has written five books, edited ten articles, and published over 300 refereed articles.

Dr. Subrahmanian is a fellow of the American Association for the Advancement of Science and the Association for the Advancement of Artificial Intelligence, and received numerous other honors and awards. His work has been featured in numerous outlets, such as the *Baltimore Sun, the Economist, Science, Nature*, the *Washington Post*, and *American Public Media*. He serves on the editorial boards of numerous journals, including *Science*, on the Board of Directors of the Development Gateway Foundation (set up by the World Bank) and SentiMetrix, Inc., and on the Research Advisory Board of Tata Consultancy Services. He previously served on DARPAs Executive Advisory Council on Advanced Logistics and as an Ad Hoc Member of the U.S. Air Force Science Advisory Board.

**Yanhai Xiong** received the bachelor's degree in automation from the University of Science and Technology of China, Hefei, China, and the Ph.D. degree in computer science and engineering from Nanyang Technological University, Singapore.

She has held a postdoctoral position at Dartmouth College, Hanover, NH, USA, since July 2018. Her research interests lie in optimization, machine learning, cybersecurity, and smart cities.