Research Statement of Albert Jiang

My research is in **multiagent systems**, an interdisciplinary field at the interface of computer science, game theory and operations research. My goal is to develop computationally efficient methods for intelligent decision-making in large-scale multiagent systems. With the ubiquity of such systems in all aspects of modern life, including electronic commerce, infrastructure security, cybersecurity, and social networks, there is great demand for better autonomous agents that help the users as well as better tools that help the system designers. Making intelligent decisions in multiagent systems requires prediction of the behavior of self-interested agents. Game-theoretic solution concepts like Nash equilibrium and correlated equilibrium are mathematical models of such self-interested behavior; however, standard computational methods fail to scale up to real-world systems.

I would characterize my research as a combination of theoretical investigations and practical applications, with close connections between the two. Methodologically, I embrace the interdisciplinary nature of my research, and have collaborated with, used techniques from, and published to diverse fields including AI, machine learning, security, theoretical computer science, operations research and economics. Throughout my PhD and postdoc I have been developing a general framework for game-theoretic computation, consisting of compact game representations, novel algorithms for solution concepts, and software implementations. I have also been applying this computational framework to real-world domains such as infrastructure security, and addressing the challenges that arise. My PhD thesis, advised by Prof. Kevin Leyton-Brown, was awarded the **Canadian Artificial Intelligence Association Doctoral Dissertation Award** and the runner up **International Foundation for Autonomous Agents and Multi-Agent Systems (IFAAMAS) Victor Lesser Distinguished Dissertation Award**. During my postdoc at Prof. Milind Tambe's TEAMCORE group at USC, I have led multiple projects that have resulted in deployments by the United States Coast Guard and the Los Angeles Sheriff's Department.

1. Research Highlights

1.1. Compact Representations for Structured Games

An important concept used throughout my research is *compact representations* of games. This is motivated by the observation that while the standard representation of games as multi-dimensional tables is impractical due to the curse of dimensionality, most real-world games have structured utility functions.

During my PhD, I led the development and analysis of **Action-Graph Games** (AGGs), a fully-expressive modeling language for representing simultaneous-move games. When games have commonly-encountered types of structured utility functions, AGGs can represent such games *compactly*, i.e., requiring small numbers of bits. Furthermore, AGGs are *computation-friendly*, that is, the structure in AGG utility functions can be exploited for efficient game-theoretic computation. In particular, I proposed a polynomial-time algorithm for the important problem of computing *expected utility* for AGGs, and leveraged this algorithm to achieve exponential speedups for finding Nash equilibria. Experiments show that using AGGs leads to a dramatic increase in the size of games accessible to computational analysis. The work is published in 2011 in Games and Economic Behavior, a top economics journal. Furthermore, in order to represent a wider variety of multiagent scenarios beyond complete-information simultaneous-move games, I extended the AGG framework to propose Temporal Action-Graph Games for representing dynamic games (UAI 2009) and Bayesian Action-Graph Games (BAGGs) for representing games of incomplete information (NIPS 2010).

My postdoc research focused on Stackelberg (leader-follower) game models of large-scale security domains, including counter-terrorism, crime suppression and environmental protection. The leader player (usually the defender) needs a randomized patrol strategy, which is computationally challenging because of the exponential number of dimensions for the strategy space. Fortunately such patrolling domains often have spatiotemporal structure. I developed a general algorithmic approach, based on a compact representation of the strategy space by fractional flows on a *transition graph*. This compact representation leads to a polynomial-sized optimization formulation which can be efficiently solved. This approach forms the basis of our applications to real-world domains including metro systems (IAAI 2012) and ferry protection (AAMAS 2013 & JAIR 2013).

The multiple types of structure in games, as well as the multiple solution concepts of interest, give rise to a diverse set of computational problems. Throughout my research experiences, I have studied and made use of a variety of techniques from AI, theoretical computer science and operations research. For example, I used tree-decomposition-based dynamic programming to compute pure-strategy Nash equilibria in AGGs (AAAI 2007). I used recent theoretical results on homomorphism problems to characterize the complexity of pure-strategy Nash equilibria for graphical games (AAMAS 2010). I used generating functions to compute pure strategy Nash equilibria in symmetric games with piecewise-linear utilities (ACM-EC 2010). I used cutting-plane methods to scale up computation of leader strategies against boundedly rational adversaries (IJCAI 2013). I used branch-and-price and submodular optimization to compute strategies for collaborating defenders (IJCAI 2013). I used derandomization and the ellipsoid method to derive a polynomial-time algorithm for the *exact* computation of correlated equilibria. This result resolved a then-open problem: previously only approximate algorithms for correlated equilibria are known. This work received the **best student paper award** at ACM-EC, the top conference on the computational aspects of economics and electronic commerce, in 2011.

1.2. From Theory to Practice

Figure 1: left, heat-map visualization of an AGG and its Nash equilibrium, generated by our software; middle, escort boats executing our strategies to protect the Staten Island Ferry; right, LASD officers using our mobile app to carry out patrols.

A major component of my PhD work consists of efforts to turn theoretical advances to practical tools that other researchers and practitioners can use. I have developed several software tools for working with AGGs and their extensions, including solvers for computing Nash equilibria for AGGs and BAGGs, a graphical user interface for creating and visualizing AGGs (Figure 1 left), and generators of AGG instances. These tools are publicly available at the AGG Project website (http://agg.cs.ubc.ca). Working with Prof. Ted Turocy, I have incorporated AGG and

BAGG functionality into the official release of GAMBIT (http://www.gambit-project.org), a widely-used software collection for game-theoretic analysis. These tools have been used by other researchers such as David Thompson to model and reason about real-world systems, including advertising auctions, network quality-of-service pricing mechanisms, and strategic voting scenarios in elections.

The TEAMCORE group has had great success with turning multiagent systems research into fielded applications with real-world impact. Since I joined as a postdoc I have participated in these efforts and gained valuable first-hand experience. I have led several projects that delivered software to security agencies. I led the TRUSTS project for the Los Angeles Sheriff's Department (LASD)'s fare enforcement effort in the LA Metro rail system. Since 2012 I also led the PROTECT project for the US Coast Guard's mission to protect ports. Another one of our PROTECT applications, for randomized scheduling of escort boats to protect ferries, has been deployed at the Staten Island Ferry in New York since April 2013 (Figure 1 middle; video at http://youtu.be/Zc5fp_L-gm4), and has been the subject of an Business Insider article. Our PROTECT team received the Coast Guard Meritorious Team Commendation in 2013. I am also involved in the writing of several grant proposals. In 2012 our proposal for randomized patrols for public transit systems was awarded $500,000 over 30 months by the TSA.

1.3. From Practice to Theory

One of the major advantages of working closely with practitioners is that through these interactions we can identify new research challenges, and resolving these challenges will in turn have real impact. For example, in the TRUSTS project for randomized fare enforcement, we initially proposed an algorithm to efficiently compute daily patrol schedules. The LASD conducted field tests of this TRUSTS system in the LA Metro in 2012, and one of the feedback comments from the officers was that patrols are often interrupted due to *execution uncertainty* such as emergencies and arrests. Utilizing techniques from planning under uncertainty, I proposed a general approach to dynamic patrolling games in uncertain environments, which provides patrol strategies with *contingency plans*. This paper was a **finalist for the best paper award** at AAMAS, a top conference on multiagent systems, in 2013. We have since developed a smart-phone app for LASD officers to execute our strategies. The LASD has conducted successful field evaluations using the app (Figure 1 right), and the TSA is currently evaluating it toward nationwide deployment.

Another important challenge arising from these applications is bounded rationality. While initial game-theoretic applications to security used the classical assumption that players are perfectly rational, in reality the adversaries are human decision makers. Recently there has been research on applying bounded-rationality models from behavior game theory to security domains, and I have done research in this aspect including human-subject experiments on Amazon Mechanical Turk for network-based security games (AAMAS 2012) and incorporating behavior models of opportunistic criminals into a game-theoretic framework for crime prevention (AAAI Fall Symposium 2013 & on-going). However, such models often have parameters that need to be estimated from data; but in many security domains data are not easily available. In a GameSec 2013 paper, I applied concepts from robust optimization, to propose a novel parameter-free solution concept called *monotonic maximin*, which provides guarantees against all adversary behavior satisfying a monotonicity condition, which includes many of the existing behavior models. I proposed an algorithm for computing monotonic maximin based on a mixed-integer linear program formulation. I am working on extending this concept beyond security domains, as a robust approach for bounded rationality in general multi-player games.

2. Research Vision

My goal is to lead a research laboratory in multiagent systems, encompassing both theoretical investigation and practical applications of game-theoretic computation. Two general themes that I anticipate will drive my work in the coming years are:

Addressing practical challenges that arise when applying game theory to the real world. I believe that game-theoretic modeling and computation can help us make more intelligent decisions in many societal challenges we face today, including cybersecurity, electronic commerce, and social networks. In particular, **cybersecurity** is becoming an increasingly important problem with the ubiquity of the Internet. I believe my expertise in security has put me in excellent position to develop practical solutions for cybersecurity. This is also a multidisciplinary domain with connections to many research areas including security, systems, and machine learning, and I am in discussion with researchers and practitioners in these relevant areas with plans to collaborate.

Developing general methodologies for efficient game-theoretic computation. In order to apply game-theoretic computation to large-scale domains like cybersecurity, we need efficient algorithms that can deal with large numbers of players as well as complex strategy spaces, such as picking paths to send data over the Internet. This type of games cannot be efficiently solved using existing technologies except in limited cases. I am developing a general algorithmic framework for such *Internet-scale multiagent systems*, extending and unifying existing approaches. I am working on a general modeling language that combines AGG-like compact utility functions and techniques from security games for compactly representing complex strategy spaces. Building on this compact representation, I plan to design efficient game-theoretic algorithms that leverage recent advances in large-scale optimization techniques form operations research. I am confident that this multi-year effort will result in an order-of-magnitude increase of the scale of domains that game theoretic analysis can be applied.