Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence (IJCAI-22)
Early Career Track

Analyzing and Designing Strategic Environments in Social Domains

Hau Chan
University of Nebraska-Lincoln
hchan3@unl.edu

Abstract

The cross-fertilization of AI and economic concepts has led to the advanced development of novel computational ideas. These ideas include models and approaches for analyzing multi-agent interaction (via game-theoretic models and solution concepts) in strategic environments and designing strategic environments (via mechanism design) to address principal decision-making problems involving multi-agent within various social contexts. In what follows, we will discuss our works on these two main topics. For analyzing multi-agent interaction, we will discuss several computational game-theoretic models to capture various agent characteristics and social (e.g., self-organization) domains. For designing strategic environments, we will discuss principal decision-making mechanism design settings in various social (e.g., facility location) contexts where the principal has to design mechanisms that elicit agent preferences over social outcomes and implement the principal's desirable social outcomes.

1 Introduction

In recent decades, we have come a long way in creating different kinds of AI that can help improve users' decision-making. Many successful AI tools are grounded in human intelligence and behavior, imitating various human decision-making components and extending human decision-making to large-scale complex multi-agent environments. Many machine learning areas (and, more notably, deep learning) are prime examples in which a fundamental idea of human learning can be used to create entire subareas that are widely applicable to many social domains. Like human interactions, complex multi-agent environments often require the consideration of agent interaction within these environments. Agents in these environments can sometimes be strategic and act noncooperatively. At the same time, these environments can sometimes be designed so that agents can act desirably and predictably under some criteria. These characteristics lead us naturally to the following research questions:

How should we model strategic environments where the agents can behave strategically in order to predict their collective behavior? How should we design strategic environments that achieve desirable agent behavior?

Much research within the AI and other CS communities, including our work, has focused on addressing these main research questions for various strategic environments in social domains using ideas that mirror human intelligence, behavior, and rationality. In particular, our work leverages the economic theories of game theory to model agent strategic behavior in strategic environments and mechanism design to design strategic environments to implement desirable agent behavior. Both game theory and mechanism design make the fundamental assumption that agents are intelligent and rational and address two very interconnected yet distinct classes of problems. The primary distinction is the knowledge of information – the strategic environment is usually presented in game theory (including key information about the agents), whereas, in mechanism design, a principal is required to design a strategic environment that elicits unknown information from the agents such that the agent behavior in the environment is desirable. Note that the elicitation component is required for the principal to analyze agent behavior in the designed environment.

The computational community's considerations of game theory and mechanism design have led to many advanced developments. More importantly, the cross-fertilization of ideas between the computation, game theory, and mechanism design communities has led to the formation of new research areas, such as algorithmic (or computational) game theory and algorithmic mechanism design. The research in these new areas facilitates the development of novel computational ideas such as models and approaches for analyzing multi-agent interaction (via game-theoretic models and solution concepts) in strategic environments and designing strategic environments (via mechanism design) to address principal decision-making problems involving multiple agents within various social contexts.

Section 2 will briefly discuss game theory and our work in computational game-theoretic models to capture various agent characteristics and social (e.g., self-organization) domains. Section 3 will briefly discuss mechanism design and our work in several principal decision-making mechanism design settings in various social (e.g., facility location) contexts.

2 Analyzing Multi-Agent Interaction

Numerous Nobel prize-winning game-theoretic models and solution concepts have been instrumental in analyzing and predicting agent behavior in various strategic social environments (or games). In a standard strategic environment with multiple agents, each agent has a set of available actions and a utility function that maps each possible combination of actions of the agent and other agents (in the environment) to some real numbers specifying the happiness of the agents. In this context, an agent's utility depends on their actions and the actions of other agents. Many of the early foundational works in economics have focused on:

1. Introducing solution or equilibrium concepts specifying which subsets of all combinations of actions of the agents as the most likely outcomes in strategic environments;
2. Predicting agent behavior (i.e., their actions) under different equilibrium concepts via mathematical characterizations in given strategic environments; and
3. Constructing game-theoretic models for specific strategic environments to study collective agent behavior.

The most well-known and well-studied solution concept is Nash equilibrium and its generalizations (e.g., correlated equilibrium), and most efforts have been devoted to studying them within strategic environments. We refer readers to existing game theory literature for more in-depth coverage of game theory.

While game theory is originally from the economic community, the AI and CS communities have significantly increased the capability and applicability of game theory by investigating computational and scalability issues surrounding computing key equilibrium concepts, representing strategic environments in large multi-agent scenarios, and applying game theory to tackle real-world issues. For instance, computing a Nash equilibrium is usually a very hard problem, and representing a large strategic environment natively can be very expensive (i.e., the representation might not fit into today's computer memory explicitly).

2.1 Research Accomplishments

Our work in the last decade has provided novel representation for strategic environments where agents have exponential numbers of actions and introduced game-theoretic models for various social domains in interdependent security investment and strategic self-organization.

Strategic Environments Where Agents have Exponential Numbers of Actions

In many real-world domains, each agent needs to make a decision that consists of multiple sub-decisions (e.g., assigning a set of resources, ranking a set of options, or finding a path in a network). As a result, each agent can have an exponential number of actions. In order to increase computational efficiency, there is a need to introduce a compact representation for representing the space of actions more compactly using less space. As the space of actions is often structured, we give a first systematic study of computation in games with structured strategy spaces. We propose Resource Graph Games (RGGs), a compact representation for games with structured strategy spaces. RGGs can compactly encode a wide range of games studied in the literature, such as congestion games, network congestion games, and various instances of security games. Not only can RGGs be used to represent existing compact classes of games, but they can also be used to model new instances of strategic environments compactly. Given the RGGs, we studied the computational aspects of computing various equilibrium concepts, including Nash equilibrium and correlated equilibrium [Chan *et al.*, 2016; Jiang *et al.*, 2017; Chan and Jiang, 2018].

Interdependent Security Investment

Over the last decades, attacks carried out by hackers and terrorists in both the physical and cyber worlds have led to increased effort by both the government and the private sector to create mechanisms that would potentially help prevent such attacks. One result of this effort has been more focused research attention to models, computational and otherwise, that facilitate the improvement of homeland security for both physical infrastructure and cyberspace. Our earlier work applied and extended game-theoretic models to study the interdependent security investment of agents, who need to determine whether to invest in security to protect against an attacker. Because there is a potential for the attack (e.g., virus and checked airline baggage) to transfer from one to other agents, the agents need to make decisions based on whether their neighbors invest in security, among other cost factors. Given our models, we studied the computational aspects of Nash equilibria to understand agent and attacker behavior and learned the parameters of the models using real-world data to conduct strategic inference and intervention [Chan *et al.*, 2012; Chan and Ortiz, 2014; Chan and Ortiz, 2015; Chan and Ortiz, 2018; Chan *et al.*, 2017; Wang *et al.*, 2019].

Strategic Self-Organization

Strategic multiagent self-organization is a natural phenomenon that occurs in various aspects of our society where strategic agents of various characteristics and (complementary or competing) objectives self-organize into groups to derive certain outcomes. Our game theory work can be characterized as modeling strategic self-organization in different social domains. These works include self-organization in crowdsourcing contest platforms [Chan *et al.*, 2020b] where contestants self-select into contests based on their skill levels and intrinsic preferences to compete for rewards, task/project contributions [Aziz *et al.*, 2021; Stevens *et al.*, 2021; Soundy *et al.*, 2021] where agents contribute to a self-selected subset of tasks or projects, residential segregation [Chan *et al.*, 2020a] where individuals self-select into living locations based on social ties and type preferences, and resource congestion [Chan and Jiang, 2016] where agents self-select subsets of resources to use depending on the resources choosing by other agents. We studied the computational aspects of Nash equilibria to understand agent self-organization behavior in these domains.

3 Designing Strategic Environments

The Nobel prize-winning economic theory of mechanism design aims to study principal decision-making problems in which a principal requires interacting with its population to implement desirable social outcomes for the population. Such interaction typically occurs in the form of information elicitation[1]: Each agent in the population has a preference over the social outcomes that is unknown to the principal, and the principal needs to elicit agent preferences to determine the outcomes according to the principal's objective, which is usually defined based on the preferences of the agents. Due to the celebrated revelation principle, the principal can design a mechanism (i.e., a strategic environment in this context) that asks the agents to reveal or report their preferences and specifies how the agent preferences will map to social outcomes. As the agents might not report their preferences accurately, many studies in mechanism design have focused on designing strategyproof mechanisms that incentivize agents to report their true preferences and output outcomes that optimize the principal's objective within specific social contexts. We refer the readers to many excellent tutorials and books for a more detailed treatment of mechanism design.

While the theory of mechanism design is originally from the economic literature, AI and, more general, CS researchers have approached mechanism design from the computational and approximability perspectives (among many others), with the usual goal of designing mechanisms that are both computationally efficient and approximately maximize the principal's objective in multiagent social domains when guaranteeing strategyproofness and efficiency/optimality simultaneously is impossible. For instance, the impossibility can occur when optimizing a principal's objective is NP-hard, or there is an incompatibility between strategyproofness and optimality.

3.1 Research Accomplishments

Our recent work in mechanism design has focused on the domain of facility location. In the most basic setting, the principal has to locate facilities (e.g., schools, parks, or libraries) within a given region to serve a set of agents. Each agent has an ideal facility location that is unknown to the principal. The principal's typical objective is to locate the facilities such that the distances of the located facilities to the agent's ideal facility locations are minimized. Since the principal does not know the agent's ideal locations, the principal aims to design strategyproof mechanisms that elicit the agent's (true) ideal locations and locate the facilities to minimize a given distance objective. Besides locating actual facilities, the facility location problem has also found a wide range of applications in other fields such as healthcare, clustering, and even problems that are not geographical, ranging from as simple as choosing the temperature for a classroom, to more advanced ones, like selecting a committee to represent people with different political views. Our recent work has introduced various variants of facility location and considered fairness in facility location.

Facilities with Capacity Constraints

Previous mechanism design studies on facility location have only focused on settings where the facilities can serve an unlimited number of agents. For public facilities such as (small) health care centers and libraries, there are often constraints (e.g., physical seating and local guidelines) limiting the maximum number of agents the facilities can serve. To capture this real-world environmental characteristic, we introduced facility location models where each facility has a capacity constraint, and each agent elects to go to the nearest facility with available capacity. Under the models, we considered strategyproof mechanisms that can approximately optimize several principal's objectives [Aziz *et al.*, 2020b; Aziz *et al.*, 2019; Aziz *et al.*, 2020a].

Agents have Ordinal Preferences over Facilities

Previous mechanism design studies on facility location have primarily focused on settings where each agent does not consider tradeoffs between their preference over the facilities and the facilities' locations. For instance, even when a facility is further away from the agent, the agent can still prefer going to that facility compared to facilities closer to the agent because the agent's underlying ordinal preference for that facility is higher. To capture the tradeoff between agent ordinal preferences and facility location distances, we introduced initial facility location models with ordinal preferences where each agent has an ordinal preference (ranking) over the facilities, and the agent's distances to the facilities are adjusted based on their ordinal preference. Under the models, we designed strategyproof mechanisms that can approximately optimize several of the principal's objectives [Chan *et al.*, 2021b].

Facility Location with Fairness Criteria

From the agent and facility perspectives, we initiated the study of fairness in facility location. From the agent perspective [Zhou *et al.*, 2022], we considered a setting where each agent is partitioned into a group based on various attributes. The principal's goal is to locate the facilities to serve groups of agents under different group-fairness objectives. Under this setting, we considered strategyproof mechanisms that can approximately optimize several principal's objectives. From the facility's perspective [Wang *et al.*, 2021], each facility has a preference over the agents it serves. The principal's goal is to allocate agents to facilities.

In addition to all of the facility location works mentioned above, in a recent IJCAI 2021 survey [Chan *et al.*, 2021a], we surveyed other existing studies and provided future directions of mechanism design for facility location.

3.2 Other Contributions in Mechanism Design

Beyond facility location, we have investigated other settings that require the principal to elicit information from the agents.

Resource Allocation with Blocking

Standard resource allocation settings studied in the literature are one-shot. However, in practice, resources are typically allocated repeatedly over time. In turn, the allocation of a given resource typically implies that such a resource will become unavailable or "blocked" for a fixed period of time. To remedy this gap between theory and practice, we propose a new repeated setting in which allocated resources can be blocked for a period of time and design (approximately) strategyproof mechanisms and algorithms for this setting [Bishop *et al.*, 2020; Bishop *et al.*, 2022].

[1]Our focus will be on mechanism design without money where the use of money to incentivize agents to report truthful information is not allowed (as in many social domains).

Budgeted Procurement

In a budgeted procurement setting, the principal has a fixed budget and wants to procure a subset of items from the sellers. Because each seller's cost is unknown to the principal, the principal has to design a (strategyproof/truthful) procurement environment or mechanism that elicits sellers' costs and selects a subset of items to procure from the sellers to maximize the principal's objective such that the total payment to the sellers is under the budget. We considered variations of the budgeted procurement setting that studies multi-unit items [Chan and Chen, 2014], a dealer's market [Chan and Chen, 2016], and proportional representation [Liu *et al.*, 2021].

4 Conclusion

The future is bright for conducting research at the intersection of AI and economic paradigms. New theoretical and practical models, solution concepts, and approaches are being developed actively to address challenges in strategic environments. There are ongoing works that continue to derive game-theoretic models and solution concepts to study scenarios in many social domains such as security (i.e., allocating limited resources to protect a set of targets) and residential segregation; both topics have appeared in AI and Economic/CS-related conferences in recent years. Many social choice and resource allocation scenarios have considered mechanism design settings to implement planners' desirable outcomes. In recent years, we have seen new practical applications of game theory and new algorithmic approaches ranging from playing Poker with humans to training machine learning models. We have also seen mechanism design applications to various types of auctions (e.g., spectrum and ads) and procurements (e.g., small businesses). As a result, there are abundant opportunities for enriching existing literature and studying new strategic environments and domains.

Acknowledgments

This work was supported by the National Institute of General Medical Sciences of the National Institutes of Health [P20GM130461] and the Rural Drug Addiction Research Center at the University of Nebraska-Lincoln. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health or the University of Nebraska.

References

[Aziz *et al.*, 2019] Haris Aziz, Hau Chan, Barton E. Lee, and David C. Parkes. The capacity constrained facility location problem. In Ioannis Caragiannis, Vahab S. Mirrokni, and Evdokia Nikolova, editors, *Web and Internet Economics*, volume 11920, page 336, 2019.

[Aziz *et al.*, 2020a] Haris Aziz, Hau Chan, Barton Lee, Bo Li, and Toby Walsh. Facility location problem with capacity constraints: Algorithmic and mechanism design perspectives. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34 of *AAAI'20*, pages 1806–1813, 2020.

[Aziz *et al.*, 2020b] Haris Aziz, Hau Chan, Barton E. Lee, and David C. Parkes. The capacity constrained facility location problem. *Games and Economic Behavior*, 124:478–490, 2020.

[Aziz *et al.*, 2021] Haris Aziz, Hau Chan, Ágnes Cseh, Bo Li, Fahimeh Ramezani, and Chenhao Wang. Multi-robot task allocation-complexity and approximation. In *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems*, AAMAS '21, pages 133–141, 2021.

[Bishop *et al.*, 2020] Nicholas Bishop, Hau Chan, Debmalya Mandal, and Long Tran-Thanh. Adversarial blocking bandits. In *Proceedings of the 34th International Conference on Neural Information Processing Systems*, NeurIPS'20, 2020.

[Bishop *et al.*, 2022] Nicholas Bishop, Hau Chan, Debmalya Mandal, and Long Tran-Thanh. Sequential blocked matching. In *Proceedings of the AAAI Conference on Artificial Intelligence*, AAAI'22, 2022.

[Chan and Chen, 2014] Hau Chan and Jing Chen. Truthful multi-unit procurements with budgets. In Tie-Yan Liu, Qi Qi, and Yinyu Ye, editors, *Web and Internet Economics*, pages 89–105, 2014.

[Chan and Chen, 2016] Hau Chan and Jing Chen. Budget feasible mechanisms for dealers. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*, AAMAS '16, pages 113–122, 2016.

[Chan and Jiang, 2016] Hau Chan and Albert Xin Jiang. Congestion games with polytopal strategy spaces. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence*, IJCAI'16, pages 165–171, 2016.

[Chan and Jiang, 2018] Hau Chan and Albert Xin Jiang. An fptas for computing nash equilibrium in resource graph games. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, IJCAI'18, pages 152–158, 2018.

[Chan and Ortiz, 2014] Hau Chan and Luis E. Ortiz. Computing nash equilibria in generalized interdependent security games. In *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 2*, NIPS'14, pages 2735–2743, 2014.

[Chan and Ortiz, 2015] Hau Chan and Luis E. Ortiz. Computing nash equilibrium in interdependent defense games. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, AAAI'15, pages 842–850, 2015.

[Chan and Ortiz, 2018] Hau Chan and Luis E. Ortiz. Learning game-theoretic models from aggregate behavioral data with applications to vaccination rates in public health. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*, AAMAS '18, pages 1891–1893, 2018.

[Chan *et al.*, 2012] Hau Chan, Michael Ceyko, and Luis E. Ortiz. Interdependent defense games: Modeling interdependent security under deliberate attacks. In *Proceedings* *of the Twenty-Eighth Conference on Uncertainty in Artificial Intelligence*, UAI'12, pages 152–162, 2012.

[Chan *et al.*, 2016] Hau Chan, Albert Xin Jiang, Kevin Leyton-Brown, and Ruta Mehta. Multilinear games. In Yang Cai and Adrian Vetta, editors, *Web and Internet Economics*, pages 44–58, 2016.

[Chan *et al.*, 2017] Hau Chan, Michael Ceyko, and Luis Ortiz. Interdependent defense games with applications to internet security at the level of autonomous systems. *Games*, 8(1), 2017.

[Chan *et al.*, 2020a] Hau Chan, Mohammad T. Irfan, and Cuong Viet Than. Schelling models with localized social influence: A game-theoretic framework. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems*, AAMAS '20, pages 240–248, 2020.

[Chan *et al.*, 2020b] Hau Chan, David C. Parkes, and Karim R. Lakhani. The price of anarchy of self-selection in tullock contests. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems*, AAMAS '20, pages 1795–1797, 2020.

[Chan *et al.*, 2021a] Hau Chan, Aris Filos-Ratsikas, Bo Li, Minming Li, and Chenhao Wang. Mechanism design for facility location problems: A survey. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence*, pages 4356–4365. IJCAI'21, 2021.

[Chan *et al.*, 2021b] Hau Chan, Minming Li, and Chenhao Wang. Facility location games with ordinal preferences. *CoRR*, abs/2107.01649, 2021.

[Jiang *et al.*, 2017] Albert Xin Jiang, Hau Chan, and Kevin Leyton-Brown. Resource graph games: A compact representation for games with structured strategy spaces. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*, AAAI'17, pages 572–578, 2017.

[Liu *et al.*, 2021] Xiang Liu, Hau Chan, Minming Li, and Weiwei Wu. Budget-feasible mechanisms for representing groups of agents proportionally. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence*, IJCAI'21, pages 313–320, 2021.

[Soundy *et al.*, 2021] Jared Soundy, Chenhao Wang, Clay Stevens, and Hau Chan. Game-theoretic analysis of effort allocation of contributors to public projects. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence*, pages 405–411. IJCAI'21, 2021.

[Stevens *et al.*, 2021] Clay Stevens, Jared Soundy, and Hau Chan. Exploring the efficiency of self-organizing software teams with game theory. In *Proceedings of the 43rd International Conference on Software Engineering: New Ideas and Emerging Results*, pages 36–40, 2021.

[Wang *et al.*, 2019] Xinrun Wang, Bo An, and Hau Chan. Who should pay the cost: A game-theoretic model for government subsidized investments to improve national cybersecurity. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, IJCAI'19, pages 6020–6027, 2019.

[Wang *et al.*, 2021] Chenhao Wang, Xiaoying Wu, Minming Li, and Hau Chan. Facility's perspective to fair facility location problems. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35 of *AAAI'21*, pages 5734–5741, 2021.

[Zhou *et al.*, 2022] Houyu Zhou, Minming Li, and Hau Chan. Strategyproof mechanisms for group-fair facility location problems. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence*. IJCAI'22, 2022.