ORMS TOMORROW

informs Student Magazine

Operations Research & Game Theory

Staff contributions:

Also included:

@ORMS_Tomorrow

facebook/ORMSTomorrow

linkedin/groups/3788040

Fall/Winter 2019

ORMS TOMORROW informs Student Magazine

Fall/Winter 2019

Dear Readers,

We are pleased to present you with the Fall/Winter 2019 edition of OR/MS Tomorrow. This issue has been made possible by the hard work of the entire staff at OR/MS Tomorrow. This thematic issue focuses on operations research and game theory. The analysis of games has been recorded as early as President James Madison's analysis on the behavior of US states under different taxation systems. Modern game theory has evolved quite a bit, with extensive developments since the 1950s. The articles in this issue serve as an introductory peak into the rich area of game theory, and its interplay with traditional operations research concepts. In addition to delving into the applications of game theory in health-care supply chains and cake cutting protocols, this issue also showcases methodological topics such as bi-level programming and machine learning. We are excited for you to read this issue. We have some exciting new developments in our team. We are pleased to welcome new members Zulqarnain Haider, Andrew Law, Feng Liu, Nithish Saji and Breanna Swan to our team. Starting February 2019, we are excited to be given the opportunity to pen down articles in a regular student column in OR/MS Today, a bi-monthly INFORMS magazine. We will also be organizing a student writing competition for the next issue of OR/MS Tomorrow (details will be announced soon!). We hope you enjoy the content and look forward to hearing any thoughts you may have regarding OR/MS Tomorrow via email at `orms_tomorrow@mail.informs.org`.

Hossein Badri and Rahul Swamy
Co-Lead Editors

Editorial Staff

A Brief Overview of Game Theory, OR, and Their Roles in Better Decision-Making

Çağlar Çağlayan
Research Associate, Robert H. Smith School of Business, University of Maryland College Park

The topic of this issue of OR/MS Tomorrow is about Game Theory and Operations Research. Accordingly, we want to give a brief overview of these two fields, introduce their fundamental concepts, and discuss their roles in better decision-making.

Operations Research (OR) is the scientific study of the management of operations and processes for the purpose of better problem-solving and decision-making (Horner, 2015). Using the tools of mathematics, statistics and computer science, OR researchers and practitioners are concerned with how managerial decisions that control the operations of the system of interest should be made and implemented to improve the targeted outcome(s) (Tanenbaum Eilon, 2018). Some of the basic concepts of OR are as follows:

- **Model:** The conceptual representation (and the mathematical formulation) of the real-life problem of interest to be simulated or mathematically solved.
- **Optimization:** Among all feasible alternatives, finding the solution with the highest achievable performance under given constraints and resources.
- **Decision Variables:** The actions/quantities to be determined from a given (possibly unbounded) set of feasible alternatives.
- **Objective Function:** The targeted outcome(s) to be optimized (i.e., maximized or minimized) by controlling the decision variables.
- **Constraints:** The equalities and inequalities that mathematically represents characteristics and the real-life limitations of the problem of interest.
- **Feasible Solution:** A specific value of decision variables satisfying all constraints.
- **Process:** The description of how the system of interest behaves/evolves, which is mathematically represented via constraints and decision variables. A process is considered "deterministic" when all parameters of the constraints are (assumed to be) known with certainty, and "stochastic" if it is inherently probabilistic and its probabilistic features are captured in the mathematical model.

Game Theory (GT) is the mathematical study of strategies, cooperation and situations involving conflicting interests, in which an agent's success in making choices (and achieving her desired outcome) depends on the choice(s) of other decision-maker(s) (Bhuiyan, 2018). It serves as a formal framework for describing social and business interactions and analyzing how decision-makers should rationally make decisions to gain the greatest possible advantage from their given situations. A few key terms that are commonly used in GT are as follows (Bhuiyan, 2018; McNulty, 2018):

- **(Strategic) Game:** The formal description of any circumstance involving a strategic interaction of two or more decision-makers with a result that depends on the decision-makers' actions. A game is called "non-cooperative" if players pursue their own interests at the expense of others' (as a result of conflicting interests).
- **Player:** Any decision-maker who takes actions in pursuit of his interests (within the context of the game) and whose actions affect the result of the game.
- **Actions:** The set of available moves that the players are allowed to do in the game, which – therefore – defines the rules of the game being played.
- **Strategy:** A complete plan of actions a player takes under particular situations in the game.
- **Payoff:** The payout/utility quantity, measuring the total satisfaction that a player receives from a particular outcome of the game.
- **Nash Equilibrium:** The stable situation, in which no player can gain any incremental benefit from changing his actions (and hence, would have no incentive to deviate from his current strategy) given the strategies of the others remain unchanged.
- **Zero-Sum Game:** A game, where no wealth is created and the net change in total utility is equal to zero as the total gains of winners are equivalent to the total losses of the other players.
- **Assumption of Rationality:** The assumption that players always make choices based on their rational outlook and strive to choose

the actions that give the outcomes they most prefer (based on their expectation on other players' strategies).

In general, the tools of OR have been used to manage organized systems and processes to achieve the optimal value for the chosen objective given the constraints and resources of the system/process of interest. Accordingly, the primary focus of the OR practitioners is on the particular system/process where the problem of interest arises, and the primary goal is to convert this real-life problem into a well-defined analytical problem that can be modeled and solved (Çağlayan, 2018). The steps of this process can be summarized as follows: (i) identifying the key decisions to be made and the key outcome (i.e., objective) to be improved, (ii) describing how these decisions affect the system behavior and the objective (via a mathematical formulation and/or programing code), (iii) using or developing the correct modeling approach capturing the constraints and key features of the system, (iv) using or designing an effective solution algorithm to identify optimal solutions, and (v) proving some of the important underlying properties of the systems of interest.

How the situation (or system) of interest is described, what its trade-offs are, and how certain actions affect the outcomes (of the situation) are also critical for GT. Yet, the focus of GT is not only on the situation itself (and its response to a decision-maker's actions) but also on what other decision-makers do and the results of their actions. This aspect of GT makes it a powerful tool for studying the situations, where the outcome cannot be predicted or assessed accurately unless the choices (of multiple decision makers) affecting the result are analyzed within the same framework (rather than in isolation).

As a final note, I would like to state my personal opinion on the role of OR and GT in decision-making. Please take it with a grain of salt as it might be "a little" biased given I am an OR person. The tools of OR such as linear programming, queuing theory and Markov decision process, are quite powerful in studying the complex systems and processes that are characterized by uncertainty, sequential decision-making, and/or many other challenging features. Accordingly, traditional OR methods might be more apt to analyze the behavior of complex systems and make better decisions for the management of such systems. Yet, the OR techniques generally only take into account what the system of interest do with respect to a decision-maker's actions rather than what other decision-makers do (and how their actions affect the situation or a decision-maker's strategy). On the other hand, GT is extremely suitable to study such circumstances and would be a more appropriate choice for analysis. As a result, depending on whether it is a complex system requiring advanced mathematical techniques to capture its key features and yield practical solutions or it is a strategic situation of cooperation or conflict where multiple decision-makers are involved, OR or GT would be the right approach offering an insightful evaluation of the particular case of study.

References

1. Horner, P. (2015, December). Meet the 'member in chief'. OR/MS Today, Vol. 42, No. 6, pp. 36-41.

2. Tanenbaum M., Eilon S., Holstein W.K., Ackoff R.L. Operations Research. In Encyclopedia Britannica online. Retrieved from https://www.britannica.com/topic/operations-research.

3. Bhuiyan, B. (2018). An Overview of Game Theory and Some Applications. Philosophy and Progress, 59(1-2), 111-128. https://doi.org/10.3329/pp.v59i1-2.36683

4. McNulty D. The Basics Of Game Theory. In Investopedia. Retrieved from https://www.investopedia.com/articles/financial-theory/08/game-theory-basics.aspixzz5XvqiB0F7

5. Çağlayan Ç. (2018, November). The Use of Quantitative Methods with Two Different Perspectives: Data-Centric versus Problem-Centric. INFORMS OR/MS Tomorrow. Retrieved from https://www.informs.org/Publications/OR-MS-Tomorrow/The-Use-of-Quantitative-Methods-with-Two-Different-Perspectives-Data-Centric-versus-Problem-Centric

Game Theory and Reinforcement Learning

Karthick Gopalswamy
Ph.D. Candidate in Industrial and Systems Engineering, North Carolina State University

This article to intended to introduce game theory and Reinforcement Learning to students in ORMS field and is not a comprehensive review of these topics. Reinforcement Learning (RL) and Game Theory (GT) are two streams of mathematics with significant applications in solving real-life problems. Despite different origins, these methods share common traits in how the problems are defined in the game environment; i.e., states, agents and strategies (or policies). Reinforcement learning, a field of machine learning, is a 'trial and error' algorithm. Based on its observations, in RL an agent acts on an unknown environment to maximize the reward. Game theory is a mathematical way of defining the logical intricacies inherent to any rational analysis of conflict. Although the terminology 'games' sounds naive, an investigation of historical games found in economics, international trade, sociology, psychology, political policies and warfare and their origins is valuable in understanding the evolution of human thinking process (evolutionary biology). Problems in RL and GT complement each other in the sense that RL provides efficient algorithms to solve more complex games using mathematical inspiration from GT. While the former is more of an art, the later could well explain the science behind this art. Both fields have a long-standing history and, despite similarities, have significantly evolved as parallel domains.

Success of Alpha Go and Open AI Five has sparked significant interest among researchers in the field of RL. RL was originally developed for solving Markov Decision Process (MDP), a stochastic process where the system is fully characterized by the given state, independent of the past. Games that are not exactly an MDP can be converted to MDP by defining the history as state; many games fall under this category. Chess is an example of a fully observable MDP where, given a state and action to be taken, the next state is known with certainty. While the underlying theory in Alpha Go was a 2-person zero-sum game for which MDP theories are well understood, this is not the case with OpenAI Five. The years between Alpha Go and OpenAI Five certainly quantify the complexity of extending 2-player games to multi-player setting. Von Neumann and Morgenstern had only managed to *define* the concept of equilibrium for a 2-person zero-sum game, a pure competition where one player gains with the loss of the other. John Nash addressed the case of competition using Nash equilibrium which, despite being a highly useful concept, uses the fundamental assumptions of rationality. This makes static Nash equilibrium, unable to extend successfully to real world dynamic problems. Nash equilibrium is defined as a set of payoff strategies with the property that no player can increase their payoff by changing their strategy alone.

Current AI systems are based on either a single agent tackling a task or a couple of agents competing (Alpha Go). Artificial General Intelligence (AGI) can materialize only by understanding how humans behave in everyday life. For example, it can be argued that the reason for helping or greeting others is because of the resulting reward; rewards like the satisfaction of helping a person and the general etiquette of greeting people. While it is easier for us to understand the result of our own actions and improve, we constantly fail to understand others. Why? We try to reason with rationality. We know people, including ourselves, are not rational. Irrationality leads to games based on imperfect information which are not susceptible of simple analysis. In my opinion, the current state of AGI is far from successful in the sense that humans tend to evolve over time and are dynamic by nature. Depending on the situation, humans choose to compete, cooperate or remain neutral. These choices increase the complexity in conceptualizing an AGI when using only the classic results on MDP, GT, etc. The classic example of prisoner's dilemma (shown in figure below) may not necessarily find equilibrium solution, when the prisoners learn based on their history.

Prisoner B / Prisoner A	Prisoner B stays silent (*cooperates*)	Prisoner B betrays (*defects*)
Prisoner A stays silent (*cooperates*)	Each serves 1 year	Prisoner A: 3 years Prisoner B: goes free
Prisoner A betrays (*defects*)	Prisoner A: goes free Prisoner B: 3 years	Each serves 2 years

Fortunately, Evolutionary Game Theory (EGT) can incorporate dynamic behavior and adaptive learning; offering a solid basis for understanding dynamic iterative situations in the context of strategic games. Further adaptation of EGT to formulate agent dynamics resulted in Multi-Agent Reinforcement Learning (MARL) which corresponds to more complex interactions. To achieve general intelligence, agents must learn how to interact with others in a shared environment: this is the challenge of MARL. Problems like negotiation among different product teams in an organization, social interactions, strategy games, and consumer markets can be addressed more accurately using MARL. With the unprecedented success of Deep Reinforcement Learning, there is a renewed interest in Deep MARL; the method used successfully by OpenAI in playing Dota-2 with human experts. Dota-2 is played in matches between two teams of five players, with each team occupying and defending their own separate base on the map. Each of the ten players independently controls a powerful character, known as a "hero", who all have unique abilities and differing styles of play. During a match, players collect experience points and items for their heroes to successfully defeat the opposing team's heroes in player versus player combat. The AI was able to learn and act in a continuously changing environment with multiple interactions (agents) over a long-time horizon. The significant leap in intelligence from Alpha Go to OpenAI Five, in my opinion, is the AI's ability to reason: strategizing with other agents (cooperate) while competing against the enemy under partially observable state (a human like behavior).

Finally, having said all the above, one must consider the moral and ethical issues surrounding the advancements of these algorithms or AGI.

While we move towards more autonomous mode of living, we should always question ourselves about accountability. AI decisions are algorithmic, unless we believe in the concept that an AI reasons like us (which is irrational), thus, we need to decide where accountability falls. On the flip side, despite not being optimal, a person will always be held accountable for the consequences of their decisions.

Saving Costs While Saving Lives

Pritha Dutta

Ph.D. Candidate in Management Science, Isenberg School of Management, University of Massachusetts at Amherst

"A new army is marching into the war against rising health care costs: engineer-mathematicians. These individuals occupy a field called operations research" wrote journalist Joel Shurkin (2013). With national discussions on healthcare topics such as the Affordable Care Act (ACA), opioid crisis, and drug pricing and shortages, it is clear that high costs and inefficient management of health care systems in the United States continue to be major concerns for government officials, policy makers and the general public.

In 2016, healthcare spending in the United States reached 3.3 trillion dollars and accounted for 17.9 percent of the country's Gross Domestic Product (GDP) according to Centers for Medicare and Medicaid Services. The figure below illustrates the widening gap between health care spending in the United States and other developed countries around the world; highlighting the need for a more efficient system in the United States.

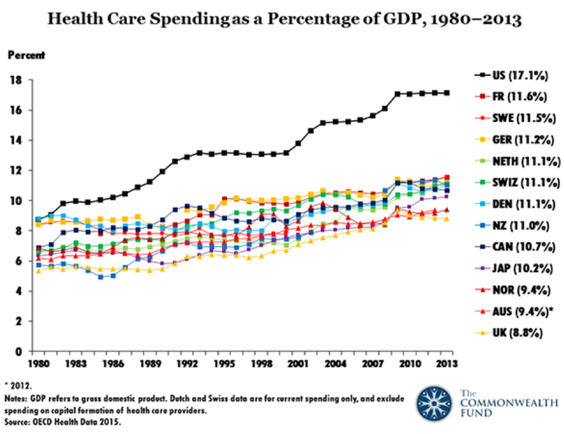

Source: Squires (2015)

It is not surprising that, being problem solvers, Operations Research scholars and practitioners are using their "weapons" to fight health care's rising costs and its immense consequences for the society at large. One such "weapon" is game theory; primarily used to analyze competition, and in many cases cooperative behavior, between stakeholders in a market or industry.

Supply chains in the health care industry are complex systems with several stakeholders such as hospitals, physicians, Group purchasing organizations (GPOs), pharmaceutical companies, medical equipment manufacturers, blood banks, insurance companies and patients who often have conflicting objectives. Game theory has been used extensively to study the economics of product supply chains. As individual stakeholders try to maximize their own gains, we end up with a system that does not function optimally. For example, if a payer (insurance company) rejects a claim made by a provider (hospital, clinic, nursing facility), the provider might benefit from providing low quality service at a lower cost. On the other hand, irrespective of the quality of service provided and cost incurred, a payer is always better off rejecting the claim (Fazulyanov, 2017). This is a classic example of the game theory concept called Prisoner's dilemma where both rational decision makers act selfishly and do not cooperate, even though it would be beneficial for both decision makers (refer to this issue's article *Game Theory and Reinforcement Learning* for more details on the Prisoner's dilemma). In the long run, this is detrimental for both the patient and the health care system. Patients may avoid seeking care if the payer rejects their claims and providers could discharge patients early to avoid costs, both leading to increased readmission rates and more severe health outcomes for the patient and, consequently, increasing costs to the system. Hence, decisions encouraging cooperation among the stakeholders in health care are critical to both reducing costs and providing the best quality of care for patients. Game theory can be used to develop strategies for negotiations between stakeholders with the ultimate goal of reducing the societal cost of health care in the United States. These strategies can be used to balance objectives between the payer and provider as described above and for interactions between other entities (provider-GPOs, GPOs-pharmaceutical companies, etc.); a few examples described in the following paragraphs.

Health care professionals often point to high prices of pharmaceutical drugs as major contributors to the high costs of health care in the United States. In addition to pricing, there are concerns and uncertainties regarding patients' response when a new drug is introduced in the market. Mahjoub, Ødegaard and Zaric (2018) used game theory to develop a pay-for-performance risk-sharing contract between a payer and a pharmaceutical company. In such a contract a new drug is prescribed to patients whose probability of response exceeds a certain threshold, then, for patients who do not respond to the drug, the manufacturer provides a rebate to the payer. Mahjoub, Ødegaard and Zaric modeled the problem as a Stackelberg game where a leader makes the first move and the follower moves sequentially. In this case the drug manufacturer is the leader who first sets the price of the drug to maximize its expected profit and the payer then decides the rebate rate and the patients who are eligible for treatment. This successful and fascinating application of game theory, among other results, found a threshold value for the rebate rate at which the net benefits for responding and nonresponding patients become equal.

Organ donation is another interesting application of the concept of the Stackelberg game. Arora and Subramanian (2016) point out the significant gap between demand and supply of organs in the United States; leading to immense socioeconomic costs. In studying an organ donation value chain (ODVC), consisting of a social planner, an organ procurement organization (OPO) and a hospital, the authors explored how the operational decisions of the OPO and the hospital affect their individual payoffs as well as social outcomes. The interactions between the (OPO) and the hospital were modeled as a Stackelberg game with the hospital acting as the leader who decides the level of effort and the operating room priority assigned to organ recovery, both of which have associated costs, while the OPO decides the level of effort to commit towards interacting with potential donors' families and seeking their authorization. In the end, administratively feasible, Pareto improving contracts were recommended to achieve optimal performance for the ODVC. Thus, showing game theory can be applied to complex problems in healthcare where efficiency in the process can save lives and ensure cost effectiveness.

Group purchasing organizations (GPOs) play a crucial role in saving costs for hospitals and health care providers. GPOs act as mediators between health care providers and manufacturers of medical supplies, blood banks, and pharmaceutical companies. Since they serve multiple health care facilities, GPOs can aggregate purchasing volumes and leverage that to get discounts from manufacturers and distributors, thereby saving costs for the health care systems. Game theory tools can be used to examine the role of GPOs in providing economies of scale in health care supply chains. Hu, Schwarz and Uhan (2012) studied how the presence of a GPO impacts the providers' total purchasing costs by developing a game-theoretic model to study a health care product supply chain consisting of a profit-maximizing manufacturer, a profit-maximizing GPO, a competitive supplier and n providers who seek to minimize their total purchasing costs. Results for this case study revealed that contract manufacturing fees charged by the GPO to the manufacturer affect the distribution of profits between manufacturers and the GPO but that contracting with the GPO did not lower the providers' total purchasing costs. This is an interesting result since the incentive for a provider to join a GPO versus directly contracting with the manufacturer is to reduce its purchasing costs.

This article displays the wide spectrum of interesting problems pertaining to health care supply chains and the various ways game theory can be applied to find a solution. Scholars in the fields of Operations Research and Management Science are studying the associated operational challenges and their socioeconomic implications by addressing issues such as nurse scheduling, organ donations, blood banking, competition among pharmaceutical firms and ambulance routing. There exists a rich body of work that uses game theory to handle some, if not all the above-mentioned problems, thus highlighting the strength of this well-established methodology in modeling complex problems, deriving optimal solutions, and potentially reducing the cost of health care delivery and services. So, the question remains, are you ready to join the "army" of OR scholars and which problem are you addressing in your next project?

References

1. Arora, P., Subramanian, R. (2016). Improving societal outcomes in the organ donation value chain. Retrieved from https://ssrn.com/abstract=2602329 or http://dx.doi.org/10.2139/ssrn.2602329.

2. Fazulyanov, I. (2017, September 18). Blockchain and Nash equilibrium in healthcare, oh my! Medium. Retrieved from https://medium.com/@ildarfazulyanov/blockchain-and-nash-equilibrium-in-healthcare-oh-my-1-6ba0e80ea03b.

3. Hu, Q., Schwarz, L. B., Uhan, N. A. (2012). The impact of group purchasing organizations on healthcare-product supply chains. Manufacturing Service Operations Management, 14(1), 7-23.

4. Mahjoub, R., Ødegaard, F., Zaric, G. S. (2018). Evaluation of a pharmaceutical risk-sharing agreement when patients are screened for the probability of success. Health economics, 27(1), 15-25.

5. Shurkin, J. (2013, April 12). Game theory tackles rising healthcare costs. Inside Science. Retrieved from https://www.insidescience.org/news/game-theory-tackles-rising-health-care-costs.

6. Squires, D. (2015, October 8). U.S. health-care from a global perspective. The Commonwealth Fund. Retrieved from https://www.commonwealthfund.org/chart/2015/health-care-spending-percentage-gdp-1980-2013.

Fair Division Problems: How to Cut Your Cake and Eat It Too

Rahul Swamy
Department of Industrial and Enterprise Systems Engineering, University of Illinois Urbana-Champaign

If there is a commodity that needs to be divided within a group of people, often times, this seemingly simple problem isn't an easy task. Each person may have a different preference on different *parts* of the commodity, one person may not know another person's preferences, the commodity may not even be divisible, and so on. This is a classical problem in game theory, aptly named *fair division*, which deals with keeping every person "happy" with their slices. From splitting rent between roommates when the rooms are of different sizes, to redrawing congressional districts for the nation's top legislative body, fair division problems are ubiquitous. Decades of research has studied and classified fair division based on the divisibility of the commodity, and the utility (or preference) distribution over the commodity.

But what exactly does 'fair' mean? A key area of research in fair division is in designing algorithms (or "protocols") that ensure certain fairness properties. It is often times important to allocate parts of a cake to people in such a way that no person would rather have another person's allocation, called *envy freeness.* In some applications, it is also prudent to ensure that the utility each person receives from their allocation is not less than the fractional weight of each person, called *proportionality.* These two criteria are mutually exclusive, and a solution to a fair division problem tries to satisfy both. Other fairness measures are also explored, such as *group-envy*, *equitability*, *exactness*, and so on. In addition to fairness, the efficiency of a division is also important. A solution is *Pareto-efficient* if it cannot be made better for one person without making it worse for another person. These notions dictate the design and analysis of fair division algorithms under suitable settings. *Cake cutting* is a fundamental part of fair division when the commodity is indivisible (i.e. can be represented in a continuous space), and the utility is heterogenous (i.e. each person values different parts of the commodity/cake differently). Cake cutting is rich in applications and mathematical theory, with many open questions to this date. For example, an interesting open question is to determine the minimum number of "cuts" required to divide a cake fairly (Magdon-Ismail et al., 2003)

How about you cut and I choose? Algorithms for cake cutting typically involves the players taking turns at cutting the cake, or one player moves a "knife" while the other players decide when to stop (Brams Taylor, 1996). The *I-cut-you-choose* algorithm is a popular protocol that works as follows for two players: Starting with the entire cake, one player divides it into two slices, and the other player chooses (freezes) the slice that they like best. This simple procedure is guaranteed to be envy-free. When the utility functions are additive functions, this also ensures proportionality. Although, when more than two players are involved, these fairness properties are not guaranteed.

Illustration of a cake cutting application to political redistricting (Spice, 2017)

A recent application of cake-cutting in a real-world setting has been in the redistricting problem. Once every ten years, the boundaries for congressional districts need to be re-drawn in order to account for migrations in populations. These districts must have roughly equal population, be geographically contiguous, and satisfy certain other legal requirements. This problem is a variation of the NP-Complete graph partitioning problem (Garey Johnson, 2002), and has been an active research topic in the ORMS literature for several decades. Recent work by Pegden et al. (2017) adapted the *I-cut-you-choose protocol* to the redistricting problem. This setting considers a game between two political parties, where at every turn, one party (A) *draws* the districts, and the other party (B) *freezes* one of the districts. In the next turn, party B draws districts from the rest of the region, and party A freezes one district, and so on. This work shows theoretical guarantees when the parties optimally "pack" and "crack" their opponent, and conditions for preserving communities of interest within the districts.

Rental harmony. A more day-to-day variation of cake-cutting is the *chore division* problem, which is also dubbed the "mirror-image of cake-cutting" since each player wants as little chore as possible. In addition to resolving conflicts among apartment room-mates, this abstract problem even has applications in dividing climate change responsibilities amongst nations (Traxler, 2002). On the algorithmic side, many cake-cutting algorithms preserve proportionality in chore-division. For two players, there is a non-trivial divide and choose algorithm called the *Selfridge-Conway protocol* which ensures envy-freeness (Robertson Webb, 1998). For an arbitrary number of players, Peterson Su (2009) was the first to provide an envy-free protocol.

As the computational world progresses towards the pursuit of fairness in automation and algorithms, the study of fair division has broadly transformable lessons in problems involving multiple conflicting stakeholders. Even though they themselves are one among several interesting game theoretic problems, the notions involved in ensuring a certain level of "happiness" among all the stakeholders is fundamental to the way we approach conflict management – at home, and outside.

References

1. Brams, S. J., Taylor, A. D. (1996). Fair Division: From cake-cutting to dispute resolution. Cambridge University Press.

2. Garey, M. R., Johnson, D. S. (2002). Computers and intractability (Vol. 29). wh freeman.

3. Magdon-Ismail, M., Busch, C., Krishnamoorthy, M. S. (2003). Cake-cutting is not a piece of cake. In Annual Symposium on Theoretical Aspects of Computer Science (pp. 596–607). Springer.

4. Pegden, W., Procaccia, A. D., Yu, D. (2017). A partisan districting protocol with provably nonpartisan outcomes. ArXiv Preprint ArXiv:1710.08781.

5. Peterson, E., Su, F. E. (2009). N-person envy-free chore division. ArXiv Preprint ArXiv:0909.0303.

6. Robertson, J., Webb, W. (1998). Cake-cutting algorithms: Be fair if you can. AK Peters/CRC Press.

7. Spice, B. (2017). "I-Cut-You-Choose" Cake-Cutting Protocol Inspires Solution to Gerrymandering. Carnegie Mellon University News. Retrieved from https://www.cmu.edu/news/stories/archives/2017/november/i-cut-you-choose-cake-cutting-protocol-inspires-solution-to-gerrymandering.html.

8. Traxler, M. (2002). Fair chore division for climate change. Social Theory and Practice, 28(1), 101–134.

Bi-level Programming with Applications in Engineering and Economics

Sepehr Ramyar
Ph.D. Candidate in Technology and Information Management, University of California Santa Cruz

Have you ever played Texas hold'em? If you have, you probably realize that the dealer has an advantage: she gets to *act* last. This allows the dealer to see what everyone else at the table is betting and then decide on what she wants to do. Having seen how everyone else has acted, the dealer can make a more informed decision. So, based on the rules of the *game*, or the sequence in which information unfolds, a player can gain an advantage over other participants. In fact, it is the *information structure* of the game (in this case, Texas hold'em) that provides an advantage for the dealer. And this is precisely why the role of the dealer rotates from one player to another, one hand at a time, so that no one player remains in a permanently advantageous position. A player does not necessarily have to be the last to act in order to be at an advantage. In fact, in many cases, being the first player to act is considered an advantage. In a duel, for example, it definitely pays to be the first player to move. Thinking of the rules of a game in terms of information structure gives us insights into how a player can be at an advantage. In the duel example, if the information (this could be how accurate they aim) is revealed *simultaneously*, a priori, no player has an advantage. However, if the information is unfolded *sequentially*, the first player to move is clearly at an advantage. In economics jargon, this is called a *first-mover advantage*.

The key notion is that the first-mover *knows* what the other player is going to do once he moves, and knowing that, he chooses an action that would optimize his objective. In other words, the advantage of the first-mover is not only in that he moves first, but also because he knows or anticipates the reaction, or the *best response*, of the other player. Although the first-mover is usually able to reduce the payoff to the other player(s) and increase his own, this is not necessarily the case and the first-mover's advantage does not always mean less payoff to the others[1] Again, the key

[1] Although it may be difficult to find real-life situations where a Stackelberg strategy would lead to higher payoffs

determinant in such situations is the information structure of the game.

But what does the information structure mean? Intuitively, this is equivalent to the first-mover *anticipating* the best response of the other player(s). In mathematical terms, one can represent a player using an optimization problem since they usually seek to maximize their payoff (or equivalently, minimize their loss). So, if a player's behavior is illustrated using a mathematical problem, the best action of a player would be the optimality condition of his optimization problem. Therefore, given what others are doing, a player's optimality condition would specify his best response to the other players. Now, if a player is including the other players' optimality conditions into his own optimization problem, he is essentially internalizing (or anticipating) the best response of other players. In other words, he is expecting how other players would respond to any of his moves and consequently impact his payoff, thus choosing the action that would maximize his payoff given the other players' anticipated reactions to his move.

This problem is called a Stackelberg game, or a leader-follower game. And the framework can be applied to many real world applications in which there is a *leader* and one (or several) *followers*. For example, a firm that can anticipate how its competitors would respond to changes in its quantity output can choose the optimal level of output knowing that the followers would react in a manner consistent with the optimality conditions of the leader. The Stackelberg formulation of a game is an instance of a more general class of optimization problems that are called *bi-level programs*. These are mathematical programs where the constraints (or the feasibility set) involve another optimization problem. This is because the leader, in a Stackelberg game, includes the optimality conditions of other players in the constraints of his optimization problem.

Bi-level programming has many applications. Resource allocation in adversarial environments is one prominent category. This situation arises when the game is played between a defender and an attacker (or a group of attackers). The defender aims to maximize security (or minimize loss) by utilizing scarce resources. For example, defending a wildlife resort against environmental crimes. This is an example of *green security games* where bilevel programming helps anticipate the adversary's strategy and accordingly allocates resources efficiently. There are other applications of Stackelberg games in urban crime prevention, traffic monitoring, toll pricing, supply chain management etc.

Despite wide applicability, bi-level programs have practical limitations. One major drawback is that bi-level programs are non-convex and hence difficult to solve. Even if the higher-level and lower-level programs are convex, the resulting bi-level program would still be non-convex because of the interdependency between the decision variables in the higher and lower level problems. Restricting the optimization problems at the two layers to linear programs would still yield an NP-hard bilevel program. There are theorems on very general conditions for existence of globally optimal solutions to bilevel programs, but none specifying an algorithm for achieving the global optimum. This is a major problem as the Stackelberg game involves more and more players (or more decision variables), the computation of a solution becomes increasingly difficult (Luo, Pang, Ralph, 1996; Ben-Aved Blair, 1990).

Another problem with Stackelberg games (or the equivalent bilevel programs) is the underlying assumption that the leader anticipates how other players would respond. In other words, there is an implicit assumption that the other players act (or respond) rationally. This, however, may not be the case in reality. One way to deal with this is using bounded rationality to model players' behavior. This includes using behavioral models that allow for mistakes (or errors) in a player choosing his best response and account for the players not always acting rationally and consistently. Introducing behavioral models that involve extra parameters exacerbates the computational difficulties of bilevel programs. Another implicit (but maybe not so innocuous) assumption behind Stackelberg games is that of a non-cooperative setting. In other words, players are assumed to be utility-maximizing individuals and there is no cooperation or side-payments of any kind. However, especially in the case of a leader and follower(s), it may not be too unrealistic to suspect formation of *coalitions* among players (Sinha, Fang, An, Keikintveld Tambe, 2018).

Despite theoretical and practical limitations, Stackelberg games, and bilevel programming in a broader sense, have provided valuable solutions to many challenges ranging from wildlife protection to prevention of price manipulations in markets. More recently, bilevel programming has been used for weighing software usefulness vs. privacy concerns and enables modeling the privacy-compromising adversaries. Advances in solving bilevel programs

to both players, there is nothing in theory that would prevent this. For example, consider the game of prisoner's dilemma. If the game is played (or the information is unfolded) sequentially, the outcome would be optimal for both players, contrary to the solution of the simultaneous game.

more efficiently would consequently have a significant impact on making our decisions more reliable and informed and our lives better and safer.

References

1. Ben-Ayed, O., Blair, C. E. (1990). Computational difficulties of bilevel linear programming. Operations Research, 38(3), 556-560.

2. Luo, Z. Q., Pang, J. S., Ralph, D. (1996). Mathematical programs with equilibrium constraints. Cambridge University Press.

3. Sinha, A., Fang, F., An, B., Kiekintveld, C., Tambe, M. (2018). Stackelberg Security Games: Looking Beyond a Decade of Success. In IJCAI (pp. 5494-5501).

Game Theory and Machine Learning

Shreya Gupta
Ph.D. Candidate in Operations Research and Industrial Engineering, University of Texas at Austin

With the onset of the New Year, it makes sense to review what has been accomplished in the past year; here, I briefly review three papers on machine learning and game theory. Many machine learning methodologies have been explored by either casting them into a game-theoretic framework or by using game theory on top of the existing machine learning frame-work. Popular examples include Generative Adversarial Nets (GANs) (Goodfellow et al., 2014) which correspond to a minimax two-player game between the generator and discriminator networks, hard margin support vector machine which can be modeled as a two-player zero-sum game (Aiolli et al., 2008), linear regression being modeled as a non-cooperative game (Ioannidis and Loiseau, 2013), and Adaboost (Freund and Schapire, 1997) which uses game theory for online learning. The variety of algorithms linking game theory and machine learning is vast and filled with fast moving trends (e.g, GANs); this review takes the road less travelled by first highlighting two new and interesting algorithms with diverse applications and finishing with a discussion on employing game theory for fairness.

First consider a Preference Learning (PL) algorithm, a machine learning methodology that involves learning preferences for previously unseen items (Polato and Aiolli, 2018). Predicting preferences is done by creating a rank ordering of pairs of preferences for every pattern. Aiolli and Polato highlight how a PL problem can be converted to "a two player zero-sum game where the row player P (the nature) picks a distribution over the whole set of training preferences (i.e., rows) aiming at minimizing the expected margin. Simultaneously, the opponent player Q (the learner) picks a distribution over the set of preference-feature pairs (i.e. columns) aiming at maximizing the expected margin (payo)." An approximation method that iteratively samples a subset of columns from a large game matrix is also proposed for the optimal strategy. In today's age of online shopping, online entertainment, online. . . everything, applications of preference learning is essential for such companies as Amazon, Netflix and Spotify to thrive.

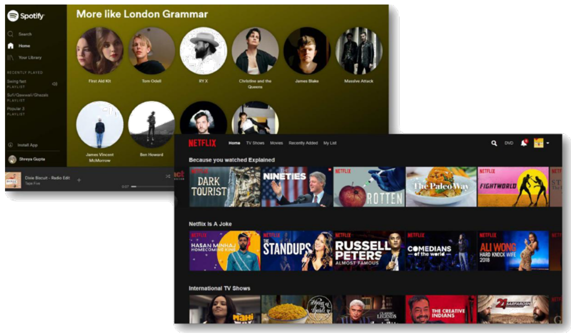

Spotify and Netflix trying to learn my preferences.

Lexical Link Analysis (LLA) is a data-driven text analysis proposed by Zhao and Zhou (2018) as a game-theoretic framework for identifying high-value information. LLA treats words as nodes and bi-grams, or word pairs, as the link between these nodes. LLA has three categories of high-value information and value metrics which are combined as the total value of information derived from the text: authoritative or popular themes, emerging themes and anomalous themes. While in social networks the popular themes are of most interest, the authors point out that emerging and anomalous themes are of higher interest in LLA due to their correlation with innovation. Deriving emerging and anomalous information from text is advantageous since these themes can develop into popular information in the future. Thus, high-value information is information that has potential to grow; this is where game theory comes in. The authors suggest a two-player framework where the information provider is one player and the rest of the world, representing the second player, responds to the information generated by the first player. In this game-theoretic approach, in addition to Nash equilibrium, another factor also important for identifying high-value information in a multi-layer network of players is that the whole system has to be Pareto efficient; i.e., the system cannot make a player better off without making another player worse off. Any application that relies heavily on text data such as user comments and descriptions in online marketplaces and hospitality services, doctor and nurse reports in electronic medical records, etc., could utilize LLA to identify high-value information. This high-value information could be discovering emerging topics of discussion in social media, proactively alleviating issues as they emerge in hospitality services, predicting product demands for out of stock or non-existent items in online marketplace inventories, and many more!

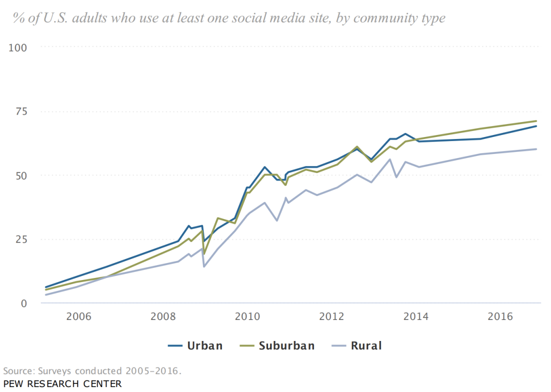

Publicly available social media content, that is made available by the informed consent of users, is a great resource for analyzing emerging trends.

Finally, I would like to discuss the relationship between fairness and game theory. First, what is fairness? The NIPS 2017 tutorial on fairness by Solon Bacrocas and Moritz Hardt (available `here`) describes fairness as “understanding and mitigating discrimination based on sensitive characteristics, such as, gender, race, religion, physical ability, and sexual orientation.” The tutorial warns that building algorithms without understanding social context, having less data for minority groups, or using a biased sample for modeling can lead to unintentional discrimination being learned by machine learning algorithms. While you can learn more about fairness from this tutorial, I want to focus your attention on the important question in Patel (2018): how do we measure and ensure an algorithm is fair? Patel (2018) helps to answer this question by exploring if fairness in algorithms should be 1) treated as a cost in the algorithm or 2) measured by an unfair prediction on the party for whom the prediction was made. Patel encourages thinking beyond the traditional goal of reduction of predictive errors in order to reduce biases embedded in data and ensure algorithmic fairness while maintaining high accuracy. To ensure algorithmic fairness, interventions can be introduced at different points in the pipeline (Friedler et al.). Preprocessing approaches assume the biases are in the training data whereas post-processing approaches work to improve interpretability and transparency so as to avoid an unfair impact on any group. Patel focuses on post-processing methods and points out that the concept of fairness in economics, revolving around the division of resources and ensuring Pareto-optimality (no one is worse off because of the decision (Corbett-Davies and Goel, 2018; Liu et al., 2018)), can be extended to machine learning. To do so, the author utilizes the Nash Welfare Product (Kaneko and Nakamura, 1979, Caragiannis et al., 2016; Venkatasubramanian and Luo; Hu and Chen, 2018), a popular concept in welfare economics which attempts to combine the utility functions of every member of the society (the subjects for whom the predictions are being made) and the utility of the institution, making the predictions into a “joint product and seeks to push all towards an equilibria” (Patel, 2018; Kaneko and Nakamura, 1979). The author justifies that Nash Welfare is seen as a halfway solution between utilitarian models (“measures the welfare of a society by the sum of the individuals’ utilities” (Stark et al., 2014)) and Rawlsian welfare models (“measures the welfare of a society by the well-being of the worst off individual (the maximin criterion)” (Rawls, 2009)). This “halfway” model helps achieve both minimizing loss in prediction accuracy and maximizing utility of the individuals and groups who are subject to the prediction.

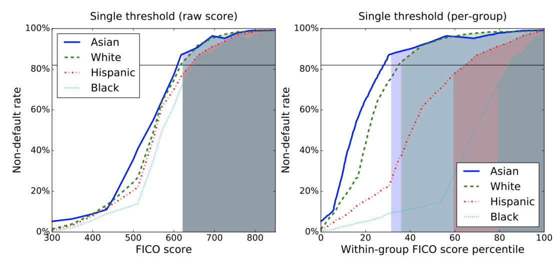

“The common FICO threshold of 620 corresponds to a non-default rate of 82%. Rescaling the x axis to represent the within-group thresholds (right), $Pr[\hat{Y} = 1|Y = 1, A]$ is the fraction of the area under the curve that is shaded. This means black non-defaulters are much less likely to qualify for loans than white or Asian ones, so a race blind score threshold violates our fairness definitions.” Hardt et al. (2016).

In conclusion, the applications of machine learning and game theory continue to expand as fields such as digital entertainment, online shopping, and service providers like hospitals seek out innovative strategies to extract customer preferences and predict new and changing trends; all while trying to ensure fairness. As the expectations of such industries evolve, it is crucial we look beyond the popular methods of the field in search for the best, most fair solutions and draw philosophically from other fields. Fairness is a concept which directly benefits from this line of thinking; the idea in Patel

(2018) of drawing from economics is an innovative attempt at combining ideas and concepts from socioeconomic sciences with that of machine learning while Bacrocas and Hardt encourage a similar approach to modeling research and development of the aforementioned tutorial. Finally, in bringing to you these ideas and papers, I might have missed out on other algorithms which also deserve to be highlighted. If there are new algorithms and modeling approaches out there that you think deserve attention, please feel free to reach out to me on `LinkedIn`.

References

1. Aiolli, F., Da San Martino, G., and Sperduti, A. (2008). A kernel method for the optimization of the margin distribution. In Kůrková, V., Neruda, R., and Koutník, J., editors, Artificial Neural Networks - ICANN 2008, pages 305–314, Berlin, Heidelberg. Springer Berlin Heidelberg.

2. Caragiannis, I., Kurokawa, D., Moulin, H., Procaccia, A. D., Shah, N., and Wang, J. (2016). The unreasonable fairness of maximum nash welfare. In Pro-ceedings of the 2016 ACM Conference on Economics and Computation, EC '16, pages 305–322, New York, NY, USA. ACM.

3. Corbett-Davies, S. and Goel, S. (2018). The measure and mismeasure of fairness: A critical review of fair machine learning. CoRR, abs/1808.00023.

4. Freund, Y. and Schapire, R. E. (1997). A decision-theoretic generalization of on-line learning and an application to boosting. Journal of Computer and System Sciences, 55(1):119 – 139.

5. Friedler, S. A., Scheidegger, C., Venkatasubramanian, S., Choudhary, S., Hamilton, E. P., and Roth, D. A comparative study of fairness-enhancing interventions in machine learning. arXiv preprint.

6. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. (2014). Generative adversarial nets. In Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N. D., and Weinberger, K. Q., editors, Advances in Neural Information Processing Systems 27, pages 2672–2680. Curran Associates, Inc.

7. Hardt, M., Price, E., and Srebro, N. (2016). Equality of opportunity in supervised learning. CoRR, abs/1610.02413.

8. Hu, L. and Chen, Y. (2018). Welfare and distributional impacts of fair classification. CoRR, abs/1807.01134.

9. Ioannidis, S. and Loiseau, P. (2013). Linear regression as a non-cooperative game. In Chen, Y. and Immorlica, N., editors, Web and Internet Eco-nomics, pages 277–290, Berlin, Heidelberg. Springer Berlin Heidelberg.

10. Kaneko, M. and Nakamura, K. (1979). The nash social wel-fare function. Econometrica, 47(2):423–435.

11. Liu, L. T., Dean, S., Rolf, E., Simchowitz, M., and Hardt, M. (2018). Delayed impact of fair machine learning. CoRR, abs/1803.04383.

12. Patel, A. (2018). Fairness for whom? critically reframing fairness with nash welfare product. CoRR, abs/1810.08540.

13. Polato, M. and Aiolli, F. (2018). Interpretable preference learning: a game theoretic framework for large margin on-line feature and rule learning. Association for the Advancement of Artificial Intelligence.

14. Rawls, J. (2009). A Theory of Justice. Harvard University Press.

15. Stark, O., Jakubek, M., and Falniowski, F. (2014). Reconciling the rawl-sian and the utilitarian approaches to the maximization of social welfare. Economics Letters, 122(3):439 – 444.

16. Venkatasubramanian, V. and Luo, Y. How much income inequality is fair? nash bargaining solution and its connection to entropy. arXiv preprint.

17. Zhao, Y. and Zhou, C. C. (2018). A game-theoretic lexical link analysis for discovering high-value information from big data. In 2018 IEEE/ACM Inter-national Conference on Advances in Social Networks Analysis and Mining (ASONAM), pages 621–625.

Applications of Game Theory in Distributed Systems: An Interview with Dr. Daniel Grosu

Hossein Badri

Ph.D. Candidate in Industrial and Systems Engineering, Wayne State University

Dr. Daniel Grosu is an Associate Professor of Computer Science and Director of the Parallel and Distributed Computing Lab at Wayne State University. Dr. Grosu and his team have published dozens of articles in top-tier peer-reviewed conferences and journals in the area of resource provisioning and pricing mechanisms in cloud, mobile edge, and vehicular edge computing systems. A main research focus of Dr. Grosu's team is the application of game theory in distributed systems, which is the subject of this interview.

Could you please briefly describe your group's research areas? We focus on a multidisciplinary research area situated at the border between computer science, game theory and economics. More specifically, we approach problems in parallel and distributed computing using techniques from game theory and economics. We focus on developing mechanisms for resource management in parallel and distributed systems, auction-based mechanisms for resource allocation in clouds and edge computing systems, mechanisms for resource allocation in competitive real-time environments, and parallel algorithms for solving non-cooperative games.

These days everybody is familiar with Cloud Computing. But what is Edge Computing? Edge Computing is a distributed computing paradigm that aims at reducing the response time of mobile applications by allowing them to perform their computation at the edge of the network instead of in cloud data centers.

What are the advantages of this new distributed computing paradigm? The main advantage of edge computing is the reduction in the response time of applications, but there are other advantages such as reliability.

From the game theory perspective, what are the main differences between cloud and edge systems? In edge systems, we have a more competitive environment for users due to the limited capacity of the edge servers compared to the capacity available from cloud data centers. We expect higher competition to acquire quality services in edge computing environments.

Going back to the game theory applications, what is the main motivation of using auction-based mechanisms in distributed systems? One of the major challenges in distributed systems, for example edge computing systems, is to decide how to allocate and price edge/cloud resources so that a given system's objective, such as revenue or social welfare, is optimized. One promising approach is to allocate these resources based on auction models, in which users place bids for using a certain amount of resources. When the auction costs are low, as is the case in the context of cloud/edge computing, auctions are especially efficient over the fixed-price markets since resources are allocated to users having the highest valuation.

Could you please give us an example of a project focused on allocating computing resources using auction-based mechanisms? One of our projects focused on designing auction-based mechanisms for resource allocation in cloud computing systems. One of the major challenges in offering infrastructure as a service (IaaS) to cloud users is designing efficient mechanisms for Virtual Machine (VM) provisioning, allocation, and pricing. Such mechanisms enable cloud providers to effectively utilize their available resources and obtain higher profits. In our setting, we allow users to request bundles of VM instances. We designed truthful greedy mechanisms for the problem such that the cloud provider provisions VMs based on the requests of the winning users and determines their payments. We showed that the proposed greedy mechanisms are truthful, that is, the users do not have incentives to manipulate the system by lying about their requested bundles of VM instances and their valuations. We also designed an incentive-compatible approximation mechanism for the problem or resource allocation and pricing in clouds. The proposed approximation mechanism drives the system into an equilibrium in which the users do not have incentives to manipulate the system by untruthfully reporting their VM bundle requests and valuations. We showed that the proposed approximation mechanism is a PTAS (Polynomial-Time Approximation Scheme) which is by far the strongest approximation result that can be achieved for this problem, unless $P = NP$.

Any examples of research projects on edge computing systems? In one of our projects, we addressed the problem of resource allocation and pricing in a two-level edge computing system. In such a system, servers with different capacities are located in the cloud or at the edge of the network. Mobile users compete for these resources and have heterogeneous demands. We designed an auction-based mechanism that allocates and prices edge/cloud resources. We showed that the proposed mechanism is individually rational and produces envy-free allocations.

What is individual-rationality and envy-freeness in this context? The individual-rationality property guarantees that users are willing to participate in the mechanism, while envy-freeness guarantees that when the auction is finished, no user would be happier with the outcome of another user.

The research projects you just described consider only one service provider. What if we have multiple providers? Considering multiple providers is very important. The amount of computing resources required by current and future data-intensive applications is expected to increase dramatically, creating high demands for cloud resources. The cloud providers' available resources may not be sufficient enough to cope with such demands. Therefore, the cloud providers need to reshape their business structures and seek to improve their dynamic resource scaling capabilities. Federated clouds offer a practical platform for addressing this service management issue. In one of our project we investigated the problem of forming federations of clouds. We introduced a cloud federation formation game that considers the cooperation of the cloud providers in offering cloud IaaS services. Based on the proposed federation formation game, we designed a cloud federation formation mechanism that enables cloud providers to dynamically form a cloud federation maximizing their profit.

Do you consider the reliability in the formation of cloud federations? Absolutely. Reliability is a key element in the formation of cloud federations. If a cloud provider agrees to provide some resources in a federation, but it fails to deliver the promised resources, then the application program could not be executed by that federation. Therefore, selecting highly trusted cloud providers to be part of the cloud federation is necessary in order to avoid this problem. In addition, a cloud provider desires to be a member of a cloud federation to obtain high profit. Therefore, a cloud federation formation mechanism should consider both profit and trust among cloud providers when making cloud federation formation decisions. One of our projects investigated the formation of reliable cloud federations.

What is the objective of the cloud federation formation problem? Usually, the objective is to maximize the individual profit of the cloud providers participating in the federations. You may also consider other objectives, such as, the overall reputation among the cloud providers.

What is your research group's plan for future research? In the future, we plan to continue our multidisciplinary research agenda working on topics at the border between computer science, game theory and economics. We plan to focus on the design of energy-aware auction-based mechanisms for application placement in edge/fog computing systems and on the design of randomized approximation mechanisms for resource allocation in such systems. We also plan to develop and investigate stochastic optimization techniques for solving resource allocation and pricing problems in edge/fog computing systems.

Student Chapter Spotlight

University at Buffalo

The UB-INFORMS Student Chapter has grown many folds over the years. Last year, the chapter broke its own record in terms of number of events and focused on establishing a strong organizational foundation. In recognition of high achievement and accelerated progress, the chapter was awarded the Summa Cum Laude award at the INFORMS 2018 Annual Meeting in Phoenix.

INFORMS 2018 Student Chapter Awards

This year, the chapter has focused on building and improving on that foundation. The chapter has covered a broad range of workshops, seminars, social and outreach events. In an effort to increase community engagement, the chapter has strengthened current relationships and built new partnerships with local organizations. To give back to the community, in addition to our pro-bono work, the chapter is currently working on running a food drive. This year has seen consistency in fundraising that has increased our spending potential with a possibility to host a variety of events.

Workshops, Seminars, and Social Events. Continuing to equip students at UB with skills that they can learn outside the classroom, this year, the chapter has hosted workshops as a knowledge sharing platform on topics selected by students. These workshops have been conducted by students who have gained expertise in these areas. A "Data Manipulation using R" workshop introduced R and its capabilities of manipulating, visualizing and analyzing data. The University at Buffalo Center for Computational Research (`CCR`) allows UB students exceptional computational resources by making super-computers accessible for research. A four-part workshop was conducted to teach students how to use this computing resource effectively. LaTeX, a critical skill, sought by many and taught by a few was one of our more popular workshops this year. The chapter hosted a two-part workshop that was open to all and received high engagement from other engineering departments.

LaTeX Workshop

In partnership with the Department of Industrial and Systems Engineering, the chapter hosted Dr. Siqian Shen from the University of Michigan, Ann Arbor, and Dr. Joyendu "Joy" Bhadury from the College at Brockport, State University of New York to present talks on their research and have dedicated meet-and-greet time with the student body. Students were also provided a platform as a part of the Distinguished Student Speaker Series to present their work and encourage research discussions. The chapter is only as good as its student body. Team-building and student networking is fostered by social events where faculty and students come together for social events, such as trivia night and darts night.

Trivia Night

Promoting Industrial and Systems Engineering. Last year, the chapter launched their first annual periodical, so-called "UB-INFORMED". This year, the chapter has extended assistance to this effort by appointing a dedicated team to lead this effort. In addition, a design competition encouraged all students to use their creative talent and develop T-shirt designs that promote Industrial and Systems Engineering.

Outreach. The UB-INFORMS team has completed a series of Pro-Bono projects. Most recently, and most notably, is the work with Western New York's Meals on Wheels (MOW). MOW is an organization that prepares and delivers meals to local seniors and disabled people and provides them with companionship. During the 2017-2018 academic year, a team of UB students partnered with MOW to create a software application that calculated the utilization of their cold storage areas. This project was completed in early October and presented to the Chief Operating Officer, Kathleen Graim and her team. The tool was well-received and they decided to move forward with a partnership. In November, the next project was outlined. An inventory management and forecasting tool is to be developed in the Spring semester.

"There is no greater feeling in business than building a product which impacts people's lives in a profound way. On behalf of Meals on Wheels for Western New York, we appreciate our partnership [with UB] and philanthropic initiatives that will increase our efficiencies and improve our service to the homebound seniors in our community." – Kathleen Graim, COO, Meals On Wheels WNY

In addition to the MOW project, another probono project team has joined forces of the UB Police. UB added a third campus in downtown Buffalo which has created a scheduling challenge for the police. This Fall, the team is creating a schedule that provides coverage to all three campuses while maintaining a work-life balance for officers.

For the students, by the students: Career Fair. This year, the chapter set a goal to expand our relationships with local businesses, and to leverage these relationships for our members. In pursuit of this goal, the chapter is currently in the process of planning a career fair, where our student body can interact with representatives from several companies in Buffalo. This career fair is additional to the career fair hosted by the school of engineering and is focused on supporting graduate students in engineering and strengthening local partnerships. As excitement about this idea grew, it was realized that such an ambitious goal might be more easily achieved through collaboration with other groups on campus. This was the catalyst for our ongoing relationship with the Graduate Indian Student Association (GISA). The UB-INFORMS chapter, specifically, is in talks with companies, such as Ford, GM, and MT Bank, among others. Tentatively planned for February 21st, 2019, we anticipate the career fair being an impactful event for a large number of graduate students at UB.

Establishing Continuity. The UB-INFORMS leadership has recognized that continuity is critical to having a successful chapter. One simple, yet impactful change, was the revision of our election process. Traditionally, e-board members and coordinators were elected during the first week of the Fall semester each year and served terms through the end of the following summer. Unfortunately, this often led to a long "start-up" time at the beginning of each academic year, as newly elected e-board members took several weeks to learn the ins and outs of their positions.

In a special meeting last April, the chapter leadership accomplished two goals. First, they formulated detailed and specific definitions for each e-board position's responsibilities, so that the constitution could act as a clear and thorough reference point in times of confusion. Second, the election process was split into two parts. Coordinator positions are still elected at the start of Fall semesters. E-board positions, however, are elected in May, three months before their term begins. This has several effects. First, it implies that e-board positions can only be held by individuals who were in the department and active in INFORMS during the preceding year. This ensures that the e-board positions are staffed by informed and engaged members of our department. Additionally, this process allows for incoming e-board members to shadow and learn from the current e-board members for several months. In this way, they are ready to hit the ground running when the new academic year arrives. Although this strategy was only implemented this year, we have already begun to see its positive effects.

With these established organizational processes and the prospective career fair, the chapter is looking forward to another successful year.

Election 2018

Student Chapter Highlights

University of Pittsburgh. The University of Pittsburgh's INFORMS chapter's mission is to prepare our graduate students for academia and industry, host talks from top researchers in the field of operations research and manufacturing and provide fun events so students can take a break from studying. This semester we hosted alumni from the banking industry to share their experiences and discuss the skills needed in industry today, organized a Python workshop covering Pandas, NumPy and Scikit for data science applications, and held a QA session with some of our department's undergraduate students to address their questions about joining grad school. To prepare our students for their talks at the INFORMS annual meeting, we held several practice sessions where students presented to each other and received feedback about their respective talks. Additionally, we held two student/faculty mixers with cake and root beer floats, a game night with the CMU INFORMS Chapter, and celebrated Thanksgiving. During this year's INFORMS Annual Meeting, our chapter was honored to receive the INFORMS 2018 Student Chapter Annual Award at the Magna Cum laude level. In the future, we are planning to hold the Annual INFORMS Trivia Night, have additional workshops and host Pitt alumni in local industry, and researchers from other schools.

University of Pittsburgh INFORMS Student Chapter

Auburn University. The INFORMS Student Chapter at Auburn University (AU) is working tirelessly to promote Operations Research and Management Science at our university. In Spring 2018, we held three great events and the election for the new board was conducted. We started the events by inviting Mr. Ventimiglia from the AU Office of the Vice President to talk about how to write proposals and how to apply for research grants for our members. We hosted a LaTeX tutorial and we visited the Kia Motors Manufacturing in Georgia. The chapter kicked off the Fall 2018 semester with collaborating on a social barbeque event on campus and continued the semester with a seminar to introduce how students can use library resources efficiently by inviting Mr. Wohrley from AU library. The students who were presenting at the 2018 INFORMS Annual Meeting had a chance to share their work with their peers in the department during the chapter's Social Presentations event. Also, the chapter provided five Travel Grants to ISE students to attend 2018 INFORMS Annual Meeting in Phoenix, AZ. In recognition of outstanding participations and performance, Auburn University Student Chapter received the 2018 "INFORMS Student Chapter Annual Award" at the level of Summa Cum Laude.

Website: `informsatau.wixsite.com/chapter`

Auburn University INFORMS Student Chapter

University of South Florida. The INFORMS student chapter at University of South Florida (INFORMS@USF) was established in 2005 with the main goal of educating USF students about different aspects of Operations Research and Management Science and expose them to the rich opportunities a career in industrial engineering can offer. Since then, our members have become actively involved in a vast variety of activities and have passed on their experience from one generation to another to maintain the chapter's excellence. Although our chapter is fairly young, in the past 10 years it has been nationally recognized with 5 Summa cum laude, 3 Magna cum laude, and 2 cum laude student chapter awards and a Judith Lieberman prize awarded at the INFORMS annual meetings. The chapter also received a Magna Cum Laude award at this year's INFORMS annual meeting in Phoenix for its activities in 2017-2018.

INFORMS@USF holds and coordinates many activities that benefit its members, the IMSE department at USF and the community at large. Some of these activities include: lecture series featuring distinguished guest researchers, social events, workshops and boot camps, community services, social activities, and community engagement with organizations and companies in the area. The chapter also coordinates with the Student Government to receive travel grants for non-presenting Masters and PhD students to provide them the much needed experience at INFORMS annual meetings. The chapter is recently collaborating with business and medical schools at USF to increase the chapter membership, and also firmly establishing its social media footprint. Furthermore, in various collaborative efforts, our members are applying OR and game theory techniques in broad areas of knowledge, including primary care, emergency care, dynamic pricing, contracts/market design, electric vehicles based shared mobility, and design of algorithms to efficiently solve OR and game theory problems.

Website: `informs.eng.usf.edu/`
Twitter: `INFORMS_USF`

INFORMS@USF

University of Massachusetts Amherst. UMass Student Chapter of INFORMS hosted multiple exciting events in the academic year 2017-2018. We started Fall 2017 with a social event where we enjoyed the lovely, fun New England tradition of apple-picking in a local orchard while also learning about the local fresh produce supply chain. As part of our INFORMS Speaker Series we hosted several talks by eminent researchers on topics from using social networks to improve driving safety to wildfire management. It was exciting to launch the chapter's YouTube Channel where interviews with our guest speakers are posted regularly. Through this initiative we aim to obtain and share insights about research topics as well life as an academic. Like every year we organized the event "Tune-ups for INFORMS Conference", where students presenting at the conference practice their talks and receive feedback from faculty as well as peers. During Spring 2018, we continued our speaker series with five more guests. One of our most popular and exciting events that we have been hosting in the Spring semester for the past two years is STEM Slam, a science communication competition open to undergraduates and graduates which has helped us in reaching out to other STEM departments across campus. It was wonderful to be recognized with the Magna Cum Laude Award at INFORMS Annual Meeting in 2017 and the Cum Laude award at this year's meeting in Phoenix for our efforts.

Website: `blogs.umass.edu/umassinf`
Twitter: `UMassINFORMS`
YouTube: `UMassINFORMS`

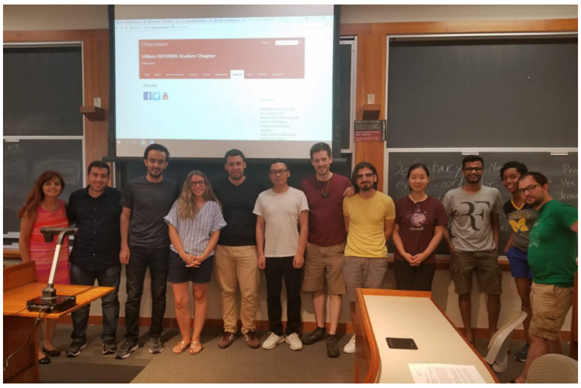

UMass INFORMS Student Chapter

Mississippi State University. The INFORMS Student chapter at Mississippi State University supports the growing interest in operations research (OR) and management science (MS) fields. We emphasize on developing academic, professional, and interpersonal skills among the student members providing a means of communication and networking among people working in OR/MS that in turn opens the window for professional and educational collaboration and enhanced opportunities. Each year we organize several information sharing sessions including seminars, professional training workshops, inviting guest speakers, and some career development events. In Fall 2018, the INFORMS Student chapter at MSU organized a social event inviting the members of other student chapters at MSU. This event was organized with a view to boost the relationships with researchers in other fields and to get acquainted with the activities hosted by other chapters. This helped us in getting some valuable insights to make our community even stronger. We also participated in INFORMS Annual Meeting 2018 at Phoenix, Arizona where we received the INFORMS 2018 student chapter award at Cum laude level. We have devoted, enthusiastic members with strong ambitions to support and extend the goal of our student chapter in the upcoming years. We also intend to enhance our collaborative efforts to boost the success of the student chapter.

Mississippi State University INFORMS Student Chapter

Do you know about the INFORMS Speaker Program?

2018 Speakers Program Committee
Gary Gaukler (chair), Sheldon H. Jacobson, and Ozlem Ergun

If your student chapter has not used the INFORMS Speakers Program before, this article is for you! Read on and find out about the fantastic benefits this program can provide for your chapter. If your chapter has used this program before, let this article serve as an inspiration to include a speaker in your upcoming events.

The Speakers Program has been in place for decades and allows INFORMS chapters (and especially student chapters) to invite INFORMS members as speakers to their chapters. The program has assembled a star-studded collection of potential speakers (all of them INFORMS members), covering topics as varied as the areas in which Operations Research, Management Sciences, and Analytics can be applied. Earlier this year, INFORMS and the Speakers Program Committee have made a few critical changes to help the program maximize its value for their members.

The focus of the program remains unchanged: connect chapters to outstanding speakers in OR, Management Science, and Analytics. Currently, there are more than 50 speakers available through the Speakers Program, with diverse topic areas including bracketology, RFID and IoT, Voting Systems, Gerrymandering, Homeland Security, Emergency Medical Services, Game Theory, and Vehicle Routing, among others. For a complete list of speakers and topics, visit www.informs.org/AvailableSpeakers. The Speakers Program continues to be an opportunity, especially for student chapters, to attract world-class speakers to their local events.

There are a number of new features of the Speakers Program. First, the Speakers Program website (www.informs.org/Resource-Center/Speakers-Program) makes it easy to browse the list of available speakers. Submitting a request for a speaker is as simple as filling in an online form with the details about the chapter requesting the speaker, the timeframe of the request, and the requested speaker(s).

Second, it has been made financially easier for chapters to host speakers through the Speakers Program. There are no honoraria to pay for the speakers; therefore, chapters are only responsible for actual hosting expenses. INFORMS recognizes that many chapters, especially the smaller chapters, lack financial resources. Therefore, INFORMS provides funding to help partially cover travel, hotel, local transportation, and/or meals. Simply submit a list of expected expenses through the online form and INFORMS will work with chapters to determine what is the budget and what can be covered. In cases of financial hardship, such as for many student chapters, INFORMS may cover up to 100% of the costs.

It is the committee's hope that many more student chapters will be motivated and encouraged to take advantage of this program. Thus, please consider the significant added value of the Speakers Program to upcoming meetings and events in your chapter and submit a speaker request at www.informs.org/RequestSpeaker.

A few notes to keep in mind if your chapter would like to use the program: All speakers are volunteers; they all have daily jobs, and hence, their availability for any specific date is subject to their own schedule. Therefore, we encourage you to request speakers at least three months in advance of the date of your potential event. It is also advisable to be flexible with either the dates and/or speakers; it is often reasonable to inquire about the availability of multiple (albeit prioritized) speakers.

We hope to see your speaker requests in the pipeline soon! If you have any further questions or suggestions, please feel free to contact Tracy Cahall at tracy.cahall@informs.org.