Operations Research

In This Issue

Operations Research
Vol. 64, No. 3, May–June 2016, pp. ii–iii
ISSN 0030-364X (print) | ISSN 1526-5463 (online)
http://dx.doi.org/10.1287/opre.2016.1518

In This Issue

Preferences Shape How We Learn from Social Networks

We all learn which decisions to make by observing the actions of our friends, a phenomenon that is called social learning. Most models of social learning in the literature assume that everyone has identical preferences. In practice, however, this is rarely true. In "Preferences, Homophily, and Social Learning," I. Lobel and E. Sadler consider a model of Bayesian learning in social networks that allow for very general network topologies, diverse preferences and even homophily (people with similar preferences are more likely to be friend). They show that density of network links is a key determinant of the impact of preferences and homophily on learning outcomes. When the network is sparse, heterogeneous preferences are harmful to learning, and homophily can lead to significant improvements. In contrast, learning outcomes are better with more diverse preferences when the network is a dense one.

A Theory to Opinions of a Society with Both Friends and Enemies

How trustful and mistrustful (antagonistic) relations shape the opinion formation of our society is an intriguing question. To provide some new insights, In "The Evolotion of Beliefs over Signed Social Networks," by G. Shi, A. Proutiiere, M. Johansson, J. Baras, and K. H. Johansson, the evolution of beliefs over signed social networks, the authors propose a model as randomized dynamics over signed networks. Two friends interact according to the DeGroot's rule. When two enemies interact, their opinions go in opposite directions. A no-survivor theorem states that the difference in opinions of any two nodes diverges whenever opinions in the entire network diverge; a live-or-die lemma indicates that almost surely the opinions either converge to an agreement or diverge. If opinions have hard lower and upper boundaries, opinions may become asymptotically clustered to, or oscillate between, the two belief boundaries, where social structural balances play a key role. A surprising fact shows that large subgroups tend to be pushed towards an opposite clustering by already clustered small subgroups.

Consensus and Information Aggregation in Symmetric Supermodular Games

In symmetric strictly supermodular games it is in players' best interest to coordinate their actions. Yet, they might fail to do so if players have incomplete and asymmetric information about the action that maximizes their payoffs. In "Learning to Coordinate in Social Networks," P. Molavi, C. Eksin, A. Ribeiro, and A. Jadbabaie provide conditions for miscoordination to eventually vanish when the game is repeatedly played. The key requirement is that the actions chosen by each player in previous plays of the game is revealed to his neighbors in a connected observation network. The information contained in these revealed actions is spread through the observation network, eventually reaching all players and allowing them to coordinate their actions. The authors also show that, under additional assumptions, the action on which players coordinate is the action that maximizes their expected payoffs given the information collectively available to them. In other words, information is efficiently aggregated through repeated plays of the game.

Coordination, Local Information, and Equilibrium Multiplicity

Many economic and social phenomena (such as bank runs, social uprisings, and the adoption of new technologies) exhibit an element of coordination, in the sense that any given agent's payoff of taking a specific action depends on the number of other agents who take the same action. In "Coordination with Local Information," M. A. Dahleh, A. Tahbaz-Salehi, J. N. Tsitsiklis, and S. I. Zoumpoulis study how local information channels shape the outcomes of coordination games. They provide conditions on the underlying information structure of the game that lead to the emergence of multiple equilibria and show that the set of equilibria shrinks as information about the fundamentals becomes more decentralized.

Competitive Equilibrium in Two Sided Matching Markets: Beyond Quasilinearity

Competitive equilibrium is a well-studied concept in matching markets. However, most of the literature on this subject deal with discrete settings or assume quasilinear utilities. Those that go beyond quasilinear utilities are usually restricted to the special case of hard budget limit, or otherwise resort to nonconstructive proofs using fixed point theorems. For arbitrary (but continuous) nonquasilinear utilities, in "Competitive Equilibria in Two Sided Matching Markets with General Utility Functions," S. Alaei, K. Jain, and Z. Malekian provide an exact characterization of utilities based on induction on the size of the market. They formulate the optimal equilibrium utility of an agent as a function of the optimal equilibrium utilities of agents on the opposite side of the market in the absence of the said agent. Subsequently, they propose a group-strategyproof mechanism that computes an equilibrium, without knowing the entire utility functions, by interactively querying agents about their indifference prices at certain points.

How Does a Network Influence Trade Delay?

It is well known that the underlying network structure of an economic system strongly influence its trade pattern. Prior literature has mainly focused on static markets. Such static models provide limited insight into how network structure impacts the speed of trade, which is an important consideration for many markets including manufacturing industries, financial and real estate markets. Why in such markets do agents sometimes fail to reach an agreement and delay trade? In "Delay in Trade Networks," T. Nguyen,

V. Subramanian, and R. Berry study a noncooperative bargaining game and show that middlemen and the complex network structure that connects them with sellers and buyers can be one cause for delay or failure in negotiation, and therefore, reduce the total trade capacity of the network. Using a connection between the bargaining game and a network circulation problem, the authors provide a sharp characterization of trade delay based on network cuts. The result gives an initial step toward network design problems for optimizing trade speed.

Optimization in Social Networks in the Face of Uncertainty

In "Online Discrete Optimization in Social Networks in the Presence of Knightian Uncertainty," M. Raginsky and A. Nedić present a model of collective real-time decision-making (or learning) in social networks in an uncertain environment, for which no a priori model is available. Learning is quantified in terms of regret, i.e., the difference between the realized performance of the network over a given time horizon and the best performance that could have been achieved in hindsight by a fictitious entity with full knowledge of the environment's evolution. Agents in the network face a time-varying sequence of composite objective functions that incorporate the total cost of their actions and the total effort they expend in deviating from their individual default strategies. The authors construct a decentralized strategy whose regret scales favorably with the time horizon, the number of agents, and the maximum number of neighbors of any agent in the social network.

Distributed Computation of Nash Equilibria over Graphs

Motivated by the proliferation of systems complicated by competition and a networked structure, the need for computing Nash equilibria over graphs is paramount. Traditionally, such a computation is carried out either in a centralized regime or in settings where the graph is fixed. In "Distributed Algorithms for aggregative games over graphs," J. Koshal, A. Nedić, and U. Shanbhag consider a subclass of Nash games (called aggregative games) being played over a graph that is evolving in time. Distributed synchronous and asynchronous schemes are presented that only necessitate local communication and represent amongst the first extensions of the consensus-based framework to a game-theoretic regime. We believe that this step is a crucial one in developing distributed protocols for a broad class of competitive decision-making problems.

Targeting over Networks

Recent advances in information technology have allowed firms to gather vast amounts of data regarding consumers' preferences and the structure and intensity of their social interactions. In "Competitive targeted Advertising over Networks," K. Bimpikis, A. Ozdaglar, and E. Yildiz explore competition among firms that target their marketing budgets to individuals who engage in word-of-mouth communication with their social peers. They provide a sharp characterization of the optimal targeted advertising strategies and highlight their dependence on the underlying network structure. Additionally, they identify settings in which firms target different subsets of individuals to soften the competition among them. Finally, they show that the firms' equilibrium advertising levels are inefficiently high and the extent of the inefficiency is increasing in the centralities of the agents they target.

Targeting the Key Leaders in Social Networks

Social networking has reshaped the human behaviors and business strategies in a fundamental manner, and various firms attempt to best utilize these social connections when spreading the new products/services among customers. In "Targeted Information Release in Social Networks" by J. Zhou and Y.-J. Chen, a theoretical model is developed to examine a pivotal step: identifying influential players when the new products/services are subject to quality uncertainty. In their setup, the network designer is allowed to release information to only a subset of players (leaders); these targeted players make their contributions first and the rest followers move subsequently after observing the leaders' decisions. Because the leaders may convey their information through their choices to the followers, the optimal selection of leaders is substantially driven by the signaling incentive. They propose a novel index for the key leader selection (i.e., a single player to provide information to), and concrete selection criteria and the pecking order for these leaders.

You Can't Hide: Finding the Source of a Rumor in a Network

Is it possible to find the source of a rumor that spreads in a network? In "Finding Rumor Sources on Random Trees," D. Shah and T. Zaman analyze this problem on different tree networks. They study the performance of an estimator known as rumor centrality, which was shown to be an exact maximum likelihood estimator for regular trees. Previous known results were very loose or restricted to very simple networks. Now, using Polya's Urn models and branching process theory, the authors have established results for a wide array of networks and rumor spreading models. They have found exact expressions for the detection and error probabilities on regular trees. They have established non-trivial detection probabilities on random expanding trees. Effectively, their results show that for any tree or tree-like network and under any rumor spreading model, the rumor source can be found, and in many cases, found with high probability.

Recommendation Algorithms for Content-Rich Applications

Modern recommender systems often use feedback from users to identify good content that can be shown to other users. This results in an explore vs. exploit tradeoff, where potentially low-quality content must be presented to users in order to learn their value. In "Online Collaborative Filtering on Graphs," S. Banerjee, S. Sanghavi, and S. Shakkottai point out that existing approaches to handling this tradeoff typically require a small and/or structured content pool (they do not however work well in content-rich settings like social network news-feeds and content-curation websites, where the amount of available content far exceeds what users can view). To this end, the authors propose an alternate framework for these settings, and present novel recommendation algorithms, with rigorous performance guarantees, under this framework. Their algorithms demonstrate the importance of serendipity in recommendation; they also suggest a potential reason behind the success of simple algorithms like Twitter's Latest-First policy.