Human Decision-Making in Multi-Agent Systems

Cao, Ming

Link to publication in University of Groningen/UMCG research database

H

Human Decision-Making in Multi-agent Systems

Ming Cao
Faculty of Science and Engineering, University of Groningen, Groningen, The Netherlands

Abstract

In order to avoid suboptimal collective behaviors and resolve social dilemmas, researchers have tried to understand how humans make decisions when interacting with other humans or smart machines and carried out theoretical and experimental studies aimed at influencing decision-making dynamics in large populations. We identify the key challenges and open issues in the related research, list a few popular models with the corresponding results, and point out future research directions.

Keywords

Decision making · Multi-agent systems

Introduction

More and more large-scale social, economic, biological, and technological complex systems have been analyzed using models of complex networks of interacting autonomous agents. Researchers are not only interested in using agent-based models to gain insight into how large dynamical networks evolve over time but also keen to introduce control actions into such networks to influence (and/or simulate) the collective behaviors of large populations. This includes social and economic networks, distributed smart energy grids, intelligent transportation systems, as well as mobile robot and smart sensor networks. However, in practice, such large numbers of interacting autonomous agents making decisions, in particular when humans are involved as participating agents, can result in highly complex, sometimes surprising, and often suboptimal, collective behaviors. It is for this very reason that human decision-making in large multi-agent networks has become a central topic for several research disciplines including economics, sociology, biology, psychology, philosophy, neuroscience, computer science, artificial intelligence, mathematics, robotics, electrical engineering, civil engineering, and last but not least systems and control.

One example to illustrate the complexity is the famous Braess paradox, observed more and more often recently in transportation, communication, and other types of flow networks, in which adding new links can actually worsen the performances of a network and vice versa (Steinberg and Zangwill 1983; Gisches and Rapoport 2012), when each agent decides to optimize its route based on its own local information. Another well-known example is the tragedy of the commons (Ostrom 2008), in which individuals use an excess of a certain shared-resource to maximize their own short-term benefit, leading to the depletion of

the resource at a great long-term cost to the group. These considerations impact on a range of emerging engineering applications, e.g., smart transportation and autonomous robots, as well as many of the most pressing current societal concerns, including strategies for energy conservation, improving recycling programs, and reducing carbon emissions.

In order to resolve these types of behavioral and social dilemmas, researchers have tried to gain insight into how humans make decisions when interacting with other autonomous agents, smart machines, or humans and carried out theoretical and experimental studies aimed at influencing decision-making dynamics in large populations in the long run. In fact, control theorists and engineers have been working at the forefront, and a collection of related works have appeared in the literature. For example, *Proceedings of the IEEE* published in 2012 a special issue on the topic of the decision-making interactions between humans and smart machines (Baillieul et al. 2012). After that, a range of new results have been reported in the literature. In what follows we first identify a few key challenges and difficulties of a number of open issues in the related research, then list a few of the most popular models and the corresponding results, and in the end point out some future research directions.

(i) Learning and adaption: Autonomous agents may learn and adapt in response to the collaboration or competition pressure from their local peers or an external changing environment.

(ii) Noise and stochastic effects: Random noise and stochastic deviation are unavoidable in both agents' decisions and the interactions among them.

(iii) Heterogeneity: Agents differ in their perception capabilities; the agents' interaction patterns co-evolve with the agents' dynamics both in space and time.

(iv) Availability for control: Some agents may not be available to be controlled directly, and even for those who are, the control is usually in the form of incentives that may only take effect in the long run or "nudge" the agent in the desired direction.

Therefore, there are still many open problems from the viewpoint of systems and control in developing a general framework for studying human decision-making in multi-agent systems; more generally, there has been an urgent need to develop new theoretical foundations together with computational and experimental tools to tackle the emerging challenging control problems associated with these evolutionary dynamics for networked autonomous agents.

Challenges and Difficulties

Significantly, when smart machines gain increasing autonomy, empowered by recent breakthroughs in data-driven cognitive learning technologies, they interact with humans in the joint decision-making processes. Humans and smart machines, collectively as networks of autonomous agents, give rise to evolutionary dynamics (Sandholm 2010) that cannot be easily modeled, analyzed, and/or controlled using current systems and control theory that has been effective thus far for engineering practices over the past decades. To highlight the evolutionary nature of the collective decision dynamics, four features can be distinguished:

Standard Models and Related Results

Models for human decision-making processes are numerous, and we list a few that are popular and becoming standard especially in the context of multi-agent systems.

Diffusion Model

The standard diffusion model was first proposed by cognitive neural scientists in the 1970s for simple sequential two-choice decision tasks and has been developed into the broader framework known as "evidence accumulation models" (Ratcliff 1978; Ratcliff et al. 2016). The internal process for a human to make a decision is taken to be a process of accumulating evidence from

some external stimulus, and once the accumulated evidence is strong enough, a decision is made. More specifically, a one-dimensional drift-diffusion process can be described by the following stochastic differential equation:

$$dz = \alpha dt + \sigma dW, \qquad z(0) = 0, \quad (1)$$

where z is the accumulated evidence in favor of a choice of interest, α is the drift rate representing the signal intensity of the stimulus acting on z, and σdW is a Wiener process with the standard deviation σ, called the diffusion rate, representing the effect of white noise. Roughly speaking, when $z(t)$ grows to reach a certain boundary level $\bar{z} > 0$ at time $\bar{t} > 0$, a decision is made at time $\bar{t}$ in favor of the choice that z corresponds to; otherwise, when another accumulator of some other choice reaches its boundary level first, then a decision favoring that choice is made. Such descriptions match the experimental data recording neural activities when human subjects make sequential decisions in lab environments (Ratcliff et al. 2016). Neural scientists have also used this model to look into the decision time and thus study the speed-accuracy trade-off in decision-making. So far, the model has been successfully applied to a range of cognitive decision-making tasks and used in clinical research (Ratcliff et al. 2016). Researchers from systems and control have looked into the convergence properties of this model (Cao et al. 2010; Woodruff et al. 2012) and used it to predict group decision-making dynamics (Stewart et al. 2012). New applications to robotic systems have also been reported (Reverdy et al. 2015).

Bayesian Model

The Bayesian model, or more generally Bayesian decision theory, is built upon the concepts of Bayesian statistics (Bernardo 1994); when making decisions, humans constantly update their estimates of the beliefs or preference values of different options using new observed inputs through Bayesian inference. It is particularly suitable for multi-alternative, multi-attribute decision-making where observations on different alternatives are dependent (Broder and Schiffer 2004; Evans et al. 2019). Let $p(A_i|x(0), x(1), \ldots, x(t))$ denote the posterior belief that alternative A_i is preferred given observations $x(j)$, $0 \le j \le t$ up to time $t \ge 0$. Then a Bayesian decision refers to the action at time t of choosing that alternative A_i among all i that maximizes the posterior probability p just given. Note that in some cases, the Bayesian model and diffusion model are closely related (Bitzer et al. 2014). The Bayesian model has found broad applications in dealing with neural and behavior data (Broder and Schiffer 2004; Evans et al. 2019).

Threshold Model

The threshold model is a classic model in sociology to study collective behavior (Granovetter 1978). It stipulates that an individual in a large population will engage in one of several possible behaviors only after a sufficiently large proportion of her surrounding individuals or the population at large have done so. Let the binary state $z_i(t)$ denote the decision of agent i in a large population at time t, $t = 0, 1, 2, \ldots$, and $\mathcal{N}_i(t)$ the set of other individuals whose decisions can be observed by agent i at time t. Then the linear threshold model dictates that

$$z_i(t+1) = \begin{cases} 1 \text{ if } \sum_{j \in \mathcal{N}_i(t)} z_j(t) \ge \Theta_i |\mathcal{N}_i(t)| \\ 0 \text{ otherwise,} \end{cases}$$

$$(2)$$

where $0 < \Theta_i < 1$ is called the *threshold* of agent i and $|\cdot|$ returns the size of the corresponding set. The model or its variants have been widely used to study propagation of beliefs and cascading of social norm violations in social networks (Centola et al. 2016; Mas and Opp 2016). Control theorists have looked into the convergence of dynamics of large populations of individuals whose behaviors follow the threshold model (Ramazi et al. 2016; Rossi et al. 2019).

Evolutionary Game Models

Game theory has been linked to decision-making processes from the day of its birth, and there are several classic textbooks covering the topic, e.g., Myerson (1991). As a branch of game theory, evolutionary game theory is

H

relevant to decision-making in large multi-agent networks especially considering the evolution of proportions of populations making certain decisions over time (Sandholm 2010). It is worth mentioning that under certain conditions, the threshold model turns out to be equivalent to some evolutionary game models (Ramazi et al. 2016). Learning strategies, such as imitation and best response, can be naturally built into different evolutionary game models; the same also holds for network structures and stochastic effects. The notions of *evolutionarily stable strategies* and *stochastically stable strategies* play prominent roles in a variety of studies (Sandholm 2010). There have been reviews on the analysis and control of evolutionary game dynamics, and we refer the interested reader to Riehl et al. (2018).

Summary and Future Directions

Human decision-making in multi-agent systems remains a fascinating and challenging topic even after various models have been proposed, tested, and compared both theoretically and experimentally in the past few decades. There are several research directions that keep gaining momentum across different disciplines:

(i) Using dynamical system models to help match human decision-making behavioral data with neural activities in the brain and provides a physiological foundation for decision theories: Although we have listed the diffusion model and Bayesian model and mentioned a few related works, the gap between behavioral and neural findings is still significant, and there is still no widely accepted unified framework that can accommodate both.

(ii) Better embedding of human decision theory into the fast developing field of network science so that the network effects in terms of information flow and adaptive learning can be better utilized to understand how a human makes decisions within a group of autonomous agents: Threshold models and evolutionary game models are developments in this direction.

(iii) Influencing and even controlling the decision-making processes: This is still a new area that requires many more novel ideas and accompanying theoretical and empirical analysis. Behavioral and social dilemmas are common in practice, and difficult to prevent for the betterment of society. Systems and control theory can play a major rule in looking for innovative decision policies and intervention rules (Riehl et al. 2018) to steer a network of autonomous agents away from suboptimal collective behaviors and toward behaviors desirable for society.

Cross-References

▶ Controlling Collective Behaviour in Complex Systems
▶ Cyber-Physical Human Systems
▶ Dynamical Social Networks
▶ Evolutionary Games
▶ Learning in Games
▶ Stochastic Games and Learning

Bibliography

Baillieul J, Leonard NE, Morgansen KA (2012) Interaction dynamics: the interface of humans and smart machines. Proc IEEE 100:567–570

Bernardo JM (1994) Bayesian theory. Wiley, Chichester

Bitzer S, Park H, Blankenburg F, Kiebel SJ (2014) Perceptual decision making: drift-diffusion model is equivalent to a Bayesian model. Front Hum Neurosci (102) 8:1–17

Broder A, Schiffer S (2004) Bayesian strategy assessment in multi-attribute decision-making. J Behav Decis Mak 16:193–213

Cao M, Stewart A, Leonard NE (2010) Convergence in human decision-making dynamics. Syst Control Lett 59:87–97

Centola D, Willer R, Macy M (2016) The emperor's dilemma: a computational model of self-enforcing norms. Am J Sociol 110(4):1009–1040

Evans NJ, Holmes WR, Trueblood JS (2019) Response-time data provide critical constraints on dynamic models of multi-alternative, multi-attribute choice. Psychon Bull Rev 26:193–213

Gisches EJ, Rapoport A (2012) Degrading network capacity may improve performance: private versus public monitoring in the Braess paradox. Theor Decis 73: 901–933

Granovetter M (1978) Threshold models of collective behavior. Am J Sociol 83:1420–1443

Mas M, Opp K-D (2016) When is ignorance bliss? Disclosing true information and cascades of norm violation in networks. Soc. Networks 47:116–129

Myerson RB (1991) Game theory: analysis of conflict. Harvard University Press, Cambridge/London

Ostrom E (2008) Tragedy of the commons. The new palgrave dictionary of economics. Palgrave Macmillan, UK, pp 3573–3576

Ramazi P, Riehl J, Cao M (2016) Networks of conforming or nonconforming individuals tend to reach satisfactory decisions. Proc Natl Acad Sci 113(46):12985–12990

Ratcliff R (1978) A theory of memory retrieval. Psychol Rev 85:59–108

Ratcliff R, Smith PL, Brown SD, McKoon G (2016) Diffusion decision model: current issues and history. Trends Cogn Sci 20(4):260–281

Reverdy P, Llhan BD, Koditschek DE (2015) A drift-diffusion model for robotic obstacle avoidance. In: Proceedings of the 2015 IEEE/RSJ international conference on intelligent robots and systems (IROS), number 15666827. IEEE

Riehl J, Ramazi P, Cao M (2018) A survey on the analysis and control of evolutionary matrix games. Ann Rev Control 45:87–106

Rossi W, Como G, Fagnai F (2019) Threshold models of cascades in large-scale networks. IEEE Trans Netw Sci Eng 6:158–172

Sandholm WH (2010) Population games and evolutionary dynamics. MIT Press, Cambridge

Steinberg R, Zangwill WI (1983) The prevalence of Braess' paradox. Transp Sci 17:301–318

Stewart A, Cao M, Nedic A, Tomlin D, Leonard AE (2012) Towards human-robot teams: model-based analysis of human decision making in two-alternative choice tasks with social feedback. Proc IEEE 100(3):751–775

Woodruff C, Vu L, Morgansen KA, Tomlin D (2012) Deterministic modeling and evaluation of decision-making dynamics in sequential two-alternative forced choice tasks. Proc IEEE 100(3):734–750

H